

APPROXIMATION OF POINTS ON LOW-DIMENSIONAL MANIFOLDS VIA RANDOM LINEAR PROJECTIONS

MARK A. IWEN MAURO MAGGIONI
DUKE UNIVERSITY, BOX 90320
DURHAM, NC 27708-0320

ABSTRACT. This paper considers the approximate reconstruction of points, $\vec{x} \in \mathbb{R}^D$, which are close to a given compact d -dimensional submanifold, \mathcal{M} , of \mathbb{R}^D using a small number of linear measurements of \vec{x} . In particular, it is shown that a number of measurements of \vec{x} which is independent of the extrinsic dimension D suffices for highly accurate reconstruction of a given \vec{x} with high probability. Furthermore, it is also proven that all vectors, \vec{x} , which are sufficiently close to \mathcal{M} can be reconstructed with uniform approximation guarantees when the number of linear measurements of \vec{x} depends logarithmically on D . Finally, the proofs of these facts are constructive: A practical algorithm for manifold-based signal recovery is presented in the process of proving the two main results mentioned above.

1. INTRODUCTION

In this paper we present a simple reconstruction technique which facilitates compressive sensing for general classes of high-dimensional signals with low intrinsic dimension. Two types of models are often considered: sparse models and low-dimensional/manifold models. The former type of model assumes that each data point has a sparse representation in terms of a (typically known) dictionary Φ , which geometrically means that data points lie on unions of a small number of planes spanned by the elements of the dictionary [25]. The latter type of model assumes that data possesses an intrinsically low-dimensional geometrical structure, for example that of a manifold (see e.g. [43, 7, 20, 24], among many others) or a union of planes (see e.g. [45, 16, 34, 29]), motivated by many applications, for example in image processing [30], computer vision [42], and pattern recognition [34].

Given the low-intrinsic dimension of these models, it is natural to ask whether a small number of linear projections (“measurements”) of a data point, together with knowledge of the low-dimensional model, suffices to encode and reconstruct a data point. In the setting of sparsity, compressed sensing [22, 40] not only says that, under suitable assumptions [25], this is indeed possible, but a convex optimization problem leads to the stable recovery of the original data point. In the setting where data lies on a low-dimensional manifold, the work of Wakin et al. [6, 46] on manifold-based signal recovery shows that low-dimensional (random) projections provide small distortion embeddings for manifolds, but leave open the question of reconstructing a data point.

Standard compressed sensing [22, 40] deals with the approximation of vectors, $\vec{x} \in \mathbb{R}^D$, which can be sparsely represented in terms of a given $D \times n$ dictionary matrix, Φ . Note that such Φ -sparse vectors can be compactly stored in a compressed form which is easy to transmit and store. Moreover, they can be recovered from their compressed representations when necessary. This compression/recovery problem has been well studied when \vec{x} is available in its entirety before compression (see, e.g., [35]). However, in situations where \vec{x} is costly to observe one may only have the ability to collect a very small set of measurements of \vec{x} to begin with, thus making standard compression techniques inapplicable (see, e.g., [2, 1] and references therein). This is the *compressed sensing regime*, where loss-less compression must occur *before* one determines which vector components or transform coefficients are actually important. Hence, the goal of standard compressed sensing becomes to design an $m \times D$ measurement matrix M , with m as small as absolutely possible, subject to the constraint that a computationally efficient reconstruction algorithm, $\mathcal{A} : \mathbb{R}^m \rightarrow \mathbb{R}^D$, exists such that $\mathcal{A}(M\vec{x}) \approx \vec{x}$ anytime $\vec{x} \in \mathbb{R}^D$ is sufficiently compressible with respect to a given dictionary matrix Φ .

More precisely, given an integer $d \ll n$ suppose that

$$\vec{x} = \Phi(\vec{f}_d + \vec{\epsilon})$$

This research was partially supported by ONR N00014-07-1-0625, The Alfred P. Sloan Foundation, DARPA N66001-11-1-4002, NSF DMS-1045153, and NSF DMS-0847388.

where $\vec{\epsilon} \in \mathbb{R}^n$ is in the row space of $\Phi \in \mathbb{R}^{D \times n}$ and

$$\vec{f}_d = \arg \min_{\vec{y} \in \mathbb{R}^n \text{ with } \|\vec{y}\|_0 \leq d} \|\vec{x} - \Phi \vec{y}\|.$$

The goal of a compressed sensing method is to approximate \vec{x} as well as possible by approximating the at most d nonzero elements of $\vec{f}_d \in \mathbb{R}^n$. Furthermore, compressed sensing techniques aim to accomplish this task using as few linear measurements of $\vec{x} \in \mathbb{R}^D$,

$$(1) \quad \langle \vec{m}_j, \vec{x} \rangle = \langle \vec{m}_j, \Phi \vec{f}_d \rangle + \langle \vec{m}_j, \Phi \vec{\epsilon} \rangle,$$

as absolutely possible.

Let $M \in \mathbb{R}^{m \times D}$ be the matrix whose j^{th} -row is the measurement vector $\vec{m}_j \in \mathbb{R}^D$ from Equation 1 above. A compressed sensing method consists of both a choice of $M \in \mathbb{R}^{m \times D}$, and a recovery algorithm, $\mathcal{A} : \mathbb{R}^m \rightarrow \mathbb{R}^D$, such that

$$(2) \quad \|\vec{x} - \mathcal{A}(M\vec{x})\|_p = \|\vec{x} - \mathcal{A}(M\Phi \vec{f}_d + M\Phi \vec{\epsilon})\|_p \leq C_{p,q} \cdot d^{\frac{1}{p} - \frac{1}{q}} \|\vec{\epsilon}\|_q$$

in fixed ℓ_p, ℓ_q norms, $1 \leq q \leq p \leq 2$, for an absolute constant $C_{p,q} \in \mathbb{R}$. Note that $M \in \mathbb{R}^{m \times D}$ forms a compressed representation of $\vec{x} \in \mathbb{R}^D$ whenever $m < D$, which is then stably inverted by \mathcal{A} . Many recovery algorithms, \mathcal{A} , have been developed for solving this problem when Φ is a square $D \times (n = D)$ orthonormal matrix, and $M\Phi$ has either restricted isometry [13] or incoherence [23, 28] properties (e.g., see [12, 10, 11, 44, 37, 38, 36, 9]). Perhaps the best such results are achieved by $(m = O(d \log(D/d))) \times D$ measurement matrices, M , whose entries are independent and identically distributed standard Gaussian random variables. These Gaussian matrices allow for near optimal compression (i.e., a near minimal size for m) while still allowing for the existence of recovery algorithms, \mathcal{A} , which achieve Equation 2 for an arbitrarily given square orthonormal matrix Φ . Furthermore, if M is Gaussian then Φ need not be known when the measurements, $M\vec{x}$, are computed: It suffices to know Φ only during reconstruction with \mathcal{A} .

One strand of work in compressed sensing has dealt with extending the results mentioned above concerning square orthonormal matrices to include settings where Φ is a more general (i.e., rectangular) $D \times n$ matrix. The first of these results extended compressed sensing to include $D \times n$ dictionaries, Φ , whose columns are all nearly pairwise orthogonal [41]. This work shares all of the advantages of the aforementioned results concerning compressed sensing when Φ is square orthonormal matrix (e.g., nearly orthogonal Φ also do not need to be known until reconstruction via \mathcal{A}) when M is a random matrix exhibiting concentration of measure properties (e.g., if M is Gaussian as above). These results were later generalized further to allow recovery along the lines of Equation 2 when Φ has columns with less limited forms of coherence and redundancy [14] (e.g., if Φ is a tight frame).

In this paper we consider a geometric generalization of standard compressed sensing results for signals which are sparsely representable with respect to a square orthonormal matrix, Φ , by focussing instead on signals which are well represented by manifold models. More specifically, herein the $D \times n$ dictionary matrix Φ utilized in standard compressed sensing models will be replaced by a piecewise linear approximation to a given submanifold of \mathbb{R}^D . To understand why this represents a generalization, note that the set of all vectors which are at least d -sparse with respect to an orthogonal matrix Φ defines a form of Grassmannian manifold consisting of $O(D^d)$ at most d -dimensional linear subspaces of \mathbb{R}^D . Hence, standard compressed sensing methods concerning square orthonormal matrices, Φ , can be viewed as dealing with a limited class of Grassmannian manifolds. In contrast, this paper allows for the approximation of signals which belong to much more general types of submanifolds of \mathbb{R}^D .

The work herein utilizes ideas introduced by Baranuik and Wakin which demonstrate the existence of simple linear operators capable of (nearly) isometrically embedding a given compact d -dimensional submanifold of \mathbb{R}^D into $\mathbb{R}^{O(d \log D)}$ without utilizing detailed knowledge regarding the submanifold's structure [6]. In some sense, this work immediately yields measurement matrices, $M \in \mathbb{R}^{m \times D}$, for manifold-based compressed sensing. However, a complete compressed sensing strategy also requires an associated reconstruction algorithm, $\mathcal{A} : \mathbb{R}^m \rightarrow \mathbb{R}^D$, capable of accurately approximating points near the given manifold in a computationally efficient fashion. Algorithms of this kind were first considered by Wakin in [46]. Therein, Wakin showed that approximating a given point, \vec{x} , near a compact d -dimensional submanifold of \mathbb{R}^D via an $O(d \log D)$ linear measurements (i.e., see Equation 1) was possible with high probability if the measurements were randomly regenerated for each new \vec{x} . Furthermore, he concluded that achieving strong reconstruction guarantees using one fixed set of linear measurements for all possible points, \vec{x} , near a given compact submanifold of \mathbb{R}^D was difficult. However, it is important to mention that the results presented in [46] were

derived independently of any particular numerical reconstruction algorithm, \mathcal{A} . As a consequence, this line of work did not result an implementable recovery algorithm with accompanying approximation guarantees.

In this paper we propose a computationally efficient reconstruction algorithm for manifold-based compressed sensing and prove accompanying approximation guarantees. In the process, we prove that a given point, \vec{x} , near a compact d -dimensional submanifold of \mathbb{R}^D can be accurately approximated using $O(d \log d)$ linear measurements with high probability when the measurements are randomly regenerated for each new \vec{x} . This improves on previous results [46] by removing all dependence on the extrinsic dimension of the submanifold, D , from the number of linear measurements required for accurate approximation. Furthermore, we provide stability guarantees for the algorithm when one fixed set of $O(d \log D)$ linear measurements are used for all possible points, \vec{x} , near a given compact submanifold of \mathbb{R}^D . Finally, an empirical evaluation of our method indicates that it also works well in practice.

Before moving on to discuss our methods and results in more detail we hasten to add that other techniques have also been proposed for manifold-based compressed sensing since the initial work of Baranuik and Wakin. Perhaps most notable among these are the statistical methods proposed by Chen et al. [17]. Chen et al. use training data from a compact d -dimensional submanifold of \mathbb{R}^D in order to estimate the manifold data's distribution via a Gaussian mixture model composed of Gaussians whose covariance matrices are all rank $O(d)$. They then use the probability density resulting from their low-rank Gaussian mixture model to approximate points on the manifold, \vec{x} , with a maximum likelihood estimator when given only linear measurements, $M\vec{x} \in \mathbb{R}^m$. In contrast, we utilize geometric and analytic techniques herein and make no attempt to estimate the statistical properties of any observed manifold data.

1.1. Methods and Results. As discussed above, the standard compressed sensing setup assumes that the signal to be approximated has a compressible representation with respect to an orthonormal basis (or frame [14], or incoherent dictionary [41]). Although this is certainly a useful setting, there are many applications where signals might be better approximated via more geometrical considerations. For example, consider the setting where the class of potential input signals varies continuously as a function of a small number of parameters (e.g., see [47, 6, 46]). In this case it makes more sense to consider the approximate reconstruction of signals, $\vec{x} \in \mathbb{R}^D$, which are close to a given compact d -dimensional submanifold, \mathcal{M} , of \mathbb{R}^D . The optimal approximation for $\vec{x} \in \mathbb{R}^D$ is then defined to be

$$\vec{x}_{\text{opt}} = \arg \min_{\vec{y} \in \mathcal{M}} \|\vec{x} - \vec{y}\|_2.$$

In effect, \vec{x}_{opt} is the best approximation to \vec{x} on \mathcal{M} . Our objective is to approximate $\vec{x}_{\text{opt}} \in \mathcal{M} \subset \mathbb{R}^D$ given only a small number of linear measurements, $M\vec{x} \in \mathbb{R}^m$, where M is an $m \times D$ measurement matrix as above. Hence, in this paper we seek to design a measurement matrix $M \in \mathbb{R}^{m \times D}$ with m as small as absolutely possible, together with a computationally efficient reconstruction algorithm $\mathcal{A} : \mathbb{R}^m \rightarrow \mathbb{R}^D$, so that $\mathcal{A}(M\vec{x}) \approx \vec{x}$ whenever $\vec{x} \in \mathbb{R}^D$ is sufficiently close to a given compact d -dimensional submanifold of \mathbb{R}^D , \mathcal{M} .¹

Note that a manifold, \mathcal{M} , is now taking the place of the dictionary matrix, $\Phi \in \mathbb{R}^{D \times n}$, in the standard compressed sensing setup discussed above. Of course, it is unreasonable to expect that we can always have an exact representation of the signal manifold at our disposal. Instead, we assume that we have a set of locally linear approximations to the given manifold which capture the local geometric structure of the manifold's tangent spaces. In fact, such piecewise linear manifold representations are exactly the type of approximations produced by existing manifold learning algorithms like LTSA [49] and Geometric Multi-Resolution Analysis [3]. Thus, we assume that the signal manifold, \mathcal{M} , is approximated by such a method at some point. However, as in standard compressed sensing methods, the manifold-based compressed sensing strategies developed below do not require that these piecewise linear manifold representations are known when the compressed measurements, $M\vec{x} \in \mathbb{R}^m$, are collected. Approximation of the signal manifold can be put off until later when signal reconstruction takes place (i.e., one does not need a piecewise linear manifold approximation until $\mathcal{A}(M\vec{x})$ is actually computed).

Although the manifold-based compressed sensing methods developed herein will work with any locally linear approximation to the given signal manifold, \mathcal{M} , we will focus on *multiscale* piecewise linear manifold approximations to \mathcal{M} in particular. As opposed to fixed-scale locally linear approximations, multiscale representations better approximate non-smooth manifolds, and manifolds contaminated with noise [15, 3]. For example, multiscale locally linear approximation is particularly beneficial for signal processing tasks involving image manifolds, which tend to be non-differentiable in many realistic settings [47]. Hence, we formulate our compressed sensing methods below with respect

¹Put another way, we require that $\mathcal{A}(M\vec{x}) \approx \vec{x}_{\text{opt}}$ which implies that $\mathcal{A}(M\vec{x}) \approx \vec{x}$ whenever $\vec{x} \approx \vec{x}_{\text{opt}}$.

to general multiscale piecewise linear manifold approximations of the type produced by Geometric Multi-Resolution Analysis (GMRA) [3].

As mentioned above, the manifold embeddings of Baranuik and Wakin [6, 46] can be considered as manifold-based compressed sensing matrices, for which however no associated recovery algorithms were explicitly defined. Indeed, the measurement matrices, $M \in \mathbb{R}^{m \times D}$, used in the manifold-based compressed sensing methods developed below are modifications of their embedding matrices. However, unlike the embedding matrices considered in [6], the measurement matrices considered herein (nearly) isometrically embed *both* the underlying signal manifold, \mathcal{M} , and the multiscale piecewise linear approximation to \mathcal{M} into \mathbb{R}^m in a way which preserves the fidelity of the embedded multiscale locally linear approximation to the embedded image of \mathcal{M} . Accomplishing this requires us to reengineer the arguments from [6] using Johnson-Lindenstrauss embedding [32] techniques similar to those utilized in [5]. The resulting measurement matrices, M , ultimately justify this complication by allowing us to develop reconstruction algorithms which work exclusively with locally linear approximations to \mathcal{M} while still preserving approximation accuracy with respect to the true manifold, \mathcal{M} .

The reconstruction algorithm, $\mathcal{A} : \mathbb{R}^m \rightarrow \mathbb{R}^D$, proposed below consists of two well-studied computational sub-routines: a method for solving approximate nearest neighbor problems (e.g., [31, 8, 4]) in a space of dimension comparable to the intrinsic dimension of the data, and a method for solving an overdetermined least squares problem (e.g., via the singular value decomposition of the associated matrix). The algorithm works by first using the compressed measurements, $M\vec{x}$, of \vec{x} to locate the best local linear approximation to \mathcal{M} at \vec{x} . This is accomplished by running a nearest neighbor algorithm on a set of “center points” from near the manifold, each of which represents a particular linear approximation to \mathcal{M} in a neighborhood of the center point. Because \mathcal{M} has low intrinsic dimension, and the center points are arranged in a multiscale hierarchy as per [3], this search can be carried out relatively quickly. To finish, the algorithm then approximates \vec{x}_{opt} , the best approximation to \vec{x} on \mathcal{M} , by solving an overdetermined least squares problem using the linear approximation to the manifold located in the first step.

In this paper we prove two compressed sensing results for the proposed reconstruction algorithm, each of which utilizes randomly generated measurement matrices, $M \in \mathbb{R}^{m \times D}$, satisfying a different set of properties. Roughly speaking, the first result indicates that $m = O(d \log(d/\delta))$ linear measurements of a given $\vec{x} \in \mathbb{R}^D$ suffice to create a compact representation, $M\vec{x} \in \mathbb{R}^m$, from which the reconstruction algorithm, \mathcal{A} , discussed above will recover an approximation to $\vec{x}_{\text{opt}} \in \mathcal{M}$ satisfying

$$\|\vec{x} - \mathcal{A}(M\vec{x})\| < C \|\vec{x} - \vec{x}_{\text{opt}}\| + \delta.$$

Here $C \in \mathbb{R}^+$ represents a fixed universal constant, $\delta \in \mathbb{R}^+$ can be freely chosen, and \mathcal{M} is the given d -dimensional submanifold of \mathbb{R}^D . This result provides what is commonly referred to as a *nonuniform recovery result*, by which we mean that the upper bound on $\|\vec{x} - \mathcal{A}(M\vec{x})\|$ holds with high probability for *each* $\vec{x} \in \mathbb{R}^D$ over the choice of random measurement matrix.

The second theorem proven below provides a type of *uniform recovery result* which holds with high probability for *all* vectors, $\vec{x} \in \mathbb{R}^D$, of a particular class. Simply put, it asserts the existence of a D -dimensional tube around the given manifold, $T \supset \mathcal{M}$, within which accurate approximation will always take place with high probability over the choice of random measurement matrix $M \in \mathbb{R}^{m \times D}$. More specifically, the second theorem says that $m = O(d \log(D/\delta))$ linear measurements of any $\vec{x} \in T \subset \mathbb{R}^D$ suffice to create a compact representation, $M\vec{x} \in \mathbb{R}^m$, from which the reconstruction algorithm discussed above, \mathcal{A} , will recover an approximation to $\vec{x}_{\text{opt}} \in \mathcal{M}$ satisfying

$$\|\vec{x} - \mathcal{A}(M\vec{x})\| < C \|\vec{x} - \vec{x}_{\text{opt}}\|_2 + \frac{C}{\sqrt{d}} \|\vec{x} - \vec{x}_{\text{opt}}\|_1 + \delta.$$

Here, as above, $C \in \mathbb{R}^+$ represents a fixed universal constant and $\delta \in \mathbb{R}^+$ can be freely chosen.

The remainder of this paper is organized as follows: In the next section we begin by fixing terminology and reviewing relevant definitions and theorems. Having established the necessary notation, we then give precise statements of the two main results proven in this paper in Section 2.2. Finally, in Section 2.3, we conclude Section 2 with a discussion of the different types of measurement matrices, $M \in \mathbb{R}^{m \times D}$, associated with each of our two main results. In Section 3 the recovery algorithm, \mathcal{A} , is presented and analyzed. In particular, the approximation error of \mathcal{A} for a given \vec{x} , $\|\vec{x} - \mathcal{A}(M\vec{x})\|$, is bounded for each of the two possible types of measurement matrices, M , considered herein. The runtime complexity of \mathcal{A} is also determined. Next, in Section 4, the number of rows, m , required for each type of measurement matrix defined in Section 2.3 is upper bounded. This formally establishes the amount of compression

possible in our manifold-based compressed sensing schemes. To finish, the compressed sensing methods developed herein are evaluated empirically in Section 5.

2. NOTATION AND SETUP

Given $n \in \mathbb{N}$ we will define $[n]$ to be the set $\{0, 1, 2, \dots, n\} \subset \mathbb{Z}$. All norms, $\|\cdot\|$, will refer to the standard Euclidean norm unless otherwise stated. We will denote an open ball of radius $\delta \in \mathbb{R}^+$ centered at $\vec{y} \in \mathbb{R}^D$ by $\mathcal{B}_\delta(\vec{y})$. Our real valued $m \times D$ measurement matrix will always be denoted by M . Furthermore, M will always be linear Johnson-Lindenstrauss embedding [32, 27, 21, 33] of a finite set $S \subset \mathbb{R}^D$ into \mathbb{R}^m .

Definition 1. Let $\epsilon \in (0, 1/2)$, and $S \subset \mathbb{R}^D$ be finite. An $m \times D$ matrix M is a linear Johnson-Lindenstrauss embedding of S into \mathbb{R}^m if

$$(1 - \epsilon)\|\vec{u} - \vec{v}\|^2 \leq \|M\vec{u} - M\vec{v}\|^2 \leq (1 + \epsilon)\|\vec{u} - \vec{v}\|^2$$

for all $\vec{u}, \vec{v} \in S$. In this case we will say that M embeds S into \mathbb{R}^m with ϵ -distortion.

The following theorem is proven by showing that an $m \times D$ matrix with randomized entries will satisfy Definition 1 for a given set $S \subset \mathbb{R}^D$ with high probability whenever m is sufficiently large (e.g., see [21]).

Theorem 1. (See [32, 21].) Let $\epsilon \in (0, 1/2)$, and $S \subset \mathbb{R}^D$ be finite. Let $m = O(\epsilon^{-2} \log |S|)$ be a natural number. Then, there exists an $m \times D$ linear Johnson-Lindenstrauss embedding of S into \mathbb{R}^m with ϵ -distortion.

For the remainder of this paper \mathcal{M} will denote a compact d -dimensional submanifold of \mathbb{R}^D with d -dimensional volume V . We will characterize results concerning any such manifold \mathcal{M} via its reach [26], denoted $\text{reach}(\mathcal{M})$, which is defined as follows: Let

$$D(\mathcal{M}) = \{\vec{x} \in \mathbb{R}^D \mid \exists \text{ a unique } \vec{y} \in \mathcal{M} \text{ with } \|\vec{x} - \vec{y}\| = d(\vec{x}, \mathcal{M})\}$$

and

$$\text{tube}_r(\mathcal{M}) = \{\vec{x} \in \mathbb{R}^D \mid d(\vec{x}, \mathcal{M}) < r\},$$

where $d(\vec{x}, \mathcal{M})$ is the standard Hausdorff distance. We then define

$$(3) \quad \text{reach}(\mathcal{M}) = \sup\{r \geq 0 \mid \text{tube}_r(\mathcal{M}) \subset D(\mathcal{M})\}.$$

Intuitively $\text{reach}(\mathcal{M})$ is the radius of the largest possible non-self-intersecting tube around \mathcal{M} . For example, if \mathcal{M} is a d -sphere of radius r , then $\text{reach}(\mathcal{M}) = r$. The reach of a manifold is particularly useful because it allows the development of concise bounds for many manifold properties of interest (e.g., curvature, self-avoidance, packing numbers, etcetera). See [26, 39, 6, 18] for more details.

Given a compact set $S \subset \mathbb{R}^D$ we define a δ -cover of S to be any finite set $S \subset \mathbb{R}^D$ with the following property:

$$\forall \vec{x} \in S, \exists \vec{y} \in S \text{ such that } \vec{x} \in \mathcal{B}_\delta(\vec{y}).$$

We will refer to a δ -cover of S , S , as *minimal* if $|S| \leq |\tilde{S}|$ for all other δ -covers of S , \tilde{S} . Hereafter, $C_\delta(S)$ will denote a minimal δ -cover of a given compact set S in \mathbb{R}^D . The following lemma, easily proven using results from [39], bounds $|C_\delta(\mathcal{M})|$ for any compact d -dimensional Riemannian manifold, \mathcal{M} , in terms of δ and $\text{reach}(\mathcal{M})$.

Lemma 1. (See [39].) Let $\mathcal{M} \subset \mathbb{R}^D$ be a compact d -dimensional Riemannian manifold with d -dimensional volume V , and suppose that $\delta \in \mathbb{R}^+$ is less than $\text{reach}(\mathcal{M})$. Then, any minimal δ -cover of \mathcal{M} , $C_\delta(\mathcal{M})$, will have

$$|C_\delta(\mathcal{M})| < \frac{V \left(\frac{d}{2} + 1\right)^{\frac{d}{2}+1}}{2^{\frac{d}{2}} \delta^d}.$$

In order to help us develop a practical recovery algorithm we will assume we have a multiscale piecewise linear approximation of \mathcal{M} of the type yielded by GMRA [3]. Let $J \in \mathbb{N}$ and $K_0, K_1, \dots, K_J \in \mathbb{N}$. For each $j \in [J]$ we assume that we have a set of affine projectors,

$$\mathbb{P}_j = \{\mathbb{P}_{j,k} : \mathbb{R}^D \rightarrow \mathbb{R}^D \mid k \in [K_j]\},$$

which approximate \mathcal{M} at scale j . More precisely, these affine projectors will collectively satisfy the three following properties:

- (1) **Affine Projections:** Every $\mathbb{P}_{j,k}$ has both an associated vector, $\vec{c}_{j,k} \in \mathbb{R}^D$, and an associated orthogonal $d \times D$ matrix, $\Phi_{j,k}$, so that

$$\mathbb{P}_{j,k}(\vec{x}) = \Phi_{j,k}^T \Phi_{j,k}(\vec{x} - \vec{c}_{j,k}) + \vec{c}_{j,k}.$$

- (2) **Dyadic Structure:** There exist two universal constants, $C_1 \in \mathbb{R}^+$ and $C_2 \in (0, 1]$, so that the following conditions are satisfied:

- (a) $K_j \leq K_{j+1}$ for all $j \in [J - 1]$.
- (b) $\|\vec{c}_{j,k_1} - \vec{c}_{j,k_2}\| > C_1 \cdot 2^{-j}$ for all $j \in [J]$ and $k_1, k_2 \in [K_j]$ with $k_1 \neq k_2$. In other words, the $\vec{c}_{j,k}$ -vectors at each scale $j \in [J]$ are well separated from one another.
- (c) For each $j \in [J] - \{0\}$ there is exactly one well defined parent function, $p_j : [K_j] \rightarrow [K_{j-1}]$, with the property that

$$\|\vec{c}_{j,k} - \vec{c}_{j-1,p_j(k)}\| < C_2 \min_{k' \in [K_{j-1}] - \{p_j(k)\}} \|\vec{c}_{j,k} - \vec{c}_{j-1,k'}\|.$$

Together these J parent functions collectively define a tree structure on the $\vec{c}_{j,k}$ -vectors. In particular, each $\vec{c}_{0,k}$ with $k \in [K_0]$ is a root node while each $\vec{c}_{j,k}$ with $k \in [K_j]$ is a leaf.

- (3) **Multiscale Approximation:** When \mathcal{M} is sufficiently smooth the affine projectors at each scale $j \in [J]$, $\{\mathbb{P}_{j,k} \mid k \in [K_j]\}$, approximate \mathcal{M} pointwise with error $O(2^{-2j})$.

- (a) There exists a constant $j_0 \in [J - 1]$ so that $\vec{c}_{j,k} \in \text{tube}_{C_1 \cdot 2^{-j-2}}(\mathcal{M})$ for all $j \in [J] - [j_0]$ and $k \in [K_j]$. Note that j_0 is a function of the constant C_1 from Property 2b. We will generally assume that a $j_0 \in [J - 1]$ satisfying this condition exists when C_1 is chosen to be as large as possible above.
- (b) For each $j \in [J]$ and $\vec{x} \in \mathbb{R}^D$ let $k_j(\vec{x}) \in [K_j]$ be such that $\vec{c}_{j,k_j(\vec{x})}$ is one of the nearest neighbors of \vec{x} in the set $\{\vec{c}_{j',k} \mid j' = j, k \in [K_j]\}$. That is, for each $j \in [J]$, let

$$k_j(\vec{x}) = \arg \min_{k \in [K_j]} \|\vec{x} - \vec{c}_{j,k}\|.$$

Then, for each $\vec{x} \in \mathcal{M}$ there exists a constant $C \in \mathbb{R}^+$ such that

$$\|\vec{x} - \mathbb{P}_{j,k_j(\vec{x})}(\vec{x})\| \leq C \cdot 2^{-2j}$$

for all $j \in [J]$. In addition, affine projectors associated with $\vec{c}_{j,k}$ -vectors that are nearly as close to any $\vec{x} \in \mathcal{M}$ as $\vec{c}_{j,k_j(\vec{x})}$ can also accurately represent \vec{x} . Hence, for each $\vec{x} \in \mathcal{M}$ there exists a constant $\tilde{C} \in \mathbb{R}^+$ such that

$$\|\vec{x} - \mathbb{P}_{j,k'}(\vec{x})\| \leq \tilde{C} \cdot 2^{-j}$$

for all $j \in [J]$ and $k' \in [K_j]$ satisfying

$$\|\vec{x} - \vec{c}_{j,k'}\| \leq 16 \cdot \max \left\{ \|\vec{x} - \vec{c}_{j,k_j(\vec{x})}\|, C_1 \cdot 2^{-j-1} \right\}.$$

Note that the affine projectors approximate \mathcal{M} more accurately as the scale $j \in [J]$ increases. The finest scale resolution is obtained when $j = J$. See [3] for details.

The remainder of this paper is devoted to analyzing the number of measurements required in order to approximately reconstruct an arbitrary point $\vec{x} \in \mathbb{R}^D$ which is nearly on a compact d -dimensional submanifold $\mathcal{M} \subset \mathbb{R}^D$. In order to yield substantive progress we must first assume some knowledge of \mathcal{M} (i.e., our manifold-based signal dictionary). Thus, we will assume below that we have a set of affine projectors, $\{\mathbb{P}_{j,k} \mid j \in [J], k \in [K_j]\}$, for \mathcal{M} as discussed above, and will primarily focus our analysis on bounding the number of measurements, m , sufficient to accurately compute $\mathbb{P}_{j,k_j(\vec{x})}(\vec{x})$ for any given input vector $\vec{x} \in \mathbb{R}^D$ and scale $j \in [J]$.

2.1. The Goal: Approximating Manifold Data via Compressive Measurements. Let $\mathbb{P} = \{\mathbb{P}_j \mid j \in [J]\}$ be a multiscale piecewise linear approximation to \mathcal{M} as discussed above. Given such a \mathbb{P} we can accurately approximate any $\vec{x} \in \mathcal{M} \subset \mathbb{R}^D$ (e.g., see Property 3b). However, herein we are primarily interested in approximating arbitrary vectors, $\vec{x} \in \mathbb{R}^D - \mathcal{M}$, as well as they can be approximated by a nearest neighbor on the manifold, $\vec{x}_{\text{opt}} \in \mathcal{M}$. As we shall see, \mathbb{P} can be utilized for this task. The following lemma demonstrates that $\mathbb{P}_{j,k_j(\vec{x})}(\vec{x})$ approximates any vector $\vec{x} \in \mathbb{R}^D$ nearly as well as $\vec{x}_{\text{opt}} \in \mathcal{M}$ does.

Lemma 2. Let $\mathcal{M} \subset \mathbb{R}^D$ be a compact d -dimensional Riemannian submanifold of \mathbb{R}^D , and $\vec{x} \in \mathbb{R}^D$. Furthermore, let $\mathbb{P}_j = \{\mathbb{P}_{j,k} \mid k \in [K_j]\}$ be a scale $j \in [J]$ GWRA approximation to \mathcal{M} . Then,

$$\|\vec{x} - \mathbb{P}_{j,k'}(\vec{x})\| \leq 17 \|\vec{x} - \vec{x}_{\text{opt}}\| + O(2^{-j})$$

for all $k' \in [K_j]$ satisfying

$$\|\vec{x} - \vec{c}_{j,k'}\| \leq 8 \cdot \max \left\{ \|\vec{x} - \vec{c}_{j,k_j(\vec{x})}\|, C_1 \cdot 2^{-j-1} \right\}.$$

Proof: Let $\delta = \max \left\{ \|\vec{x} - \vec{c}_{j,k_j(\vec{x})}\|, C_1 \cdot 2^{-j-1} \right\}$, where $C_1 \in \mathbb{R}$ is defined as in Property 2b above. Furthermore, let $k' \in [K_j]$ be such that $\|\vec{x} - \vec{c}_{j,k'}\| \leq 8\delta$. To begin, suppose that $\|\vec{x} - \vec{c}_{j,k'}\| \leq 17 \|\vec{x} - \vec{x}_{\text{opt}}\|$. In this case we are essentially finished since

$$\|\vec{x} - \mathbb{P}_{j,k'}(\vec{x})\| = \left\| \left[I - \Phi_{j,k'}^T \Phi_{j,k'} \right] (\vec{x} - \vec{c}_{j,k'}) \right\| \leq \|\vec{x} - \vec{c}_{j,k'}\| \leq 17 \|\vec{x} - \vec{x}_{\text{opt}}\|.$$

Thus, we will hereafter assume that $\|\vec{x} - \vec{c}_{j,k'}\| > 17 \|\vec{x} - \vec{x}_{\text{opt}}\|$ without loss of generality.

Repeatedly applying the triangle inequality we see that $\|\vec{x} - \mathbb{P}_{j,k'}(\vec{x})\|$ is bounded above by

$$\|\vec{x} - \vec{x}_{\text{opt}}\| + \|\vec{x}_{\text{opt}} - \mathbb{P}_{j,k'}(\vec{x}_{\text{opt}})\| + \|\mathbb{P}_{j,k'}(\vec{x}_{\text{opt}}) - \mathbb{P}_{j,k'}(\vec{x})\|.$$

The third term in the sum immediately above can be bounded by

$$\|\mathbb{P}_{j,k'}(\vec{x}_{\text{opt}}) - \mathbb{P}_{j,k'}(\vec{x})\| = \|\Phi_{j,k'}^T \Phi_{j,k'}(\vec{x} - \vec{x}_{\text{opt}})\| \leq \|\vec{x} - \vec{x}_{\text{opt}}\|.$$

To bound the second term we note that $\|\vec{x} - \vec{c}_{j,k'}\| > 17 \|\vec{x} - \vec{x}_{\text{opt}}\|$ implies that $\|\vec{x}_{\text{opt}} - \vec{c}_{j,k_j(\vec{x}_{\text{opt}})}\| > 9 \|\vec{x} - \vec{x}_{\text{opt}}\| / 8$. Therefore,

$$\begin{aligned} \|\vec{x}_{\text{opt}} - \vec{c}_{j,k'}\| &\leq \|\vec{x} - \vec{x}_{\text{opt}}\| + \|\vec{x} - \vec{c}_{j,k'}\| \leq \|\vec{x} - \vec{x}_{\text{opt}}\| + 8 \cdot \max \left\{ \|\vec{x} - \vec{c}_{j,k_j(\vec{x}_{\text{opt}})}\|, C_1 \cdot 2^{-j-1} \right\} \\ &\leq 9 \|\vec{x} - \vec{x}_{\text{opt}}\| + 8 \cdot \max \left\{ \|\vec{x}_{\text{opt}} - \vec{c}_{j,k_j(\vec{x}_{\text{opt}})}\|, C_1 \cdot 2^{-j-1} \right\} < 16 \cdot \max \left\{ \|\vec{x}_{\text{opt}} - \vec{c}_{j,k_j(\vec{x}_{\text{opt}})}\|, C_1 \cdot 2^{-j-1} \right\}. \end{aligned}$$

Property 3b now guarantees that $\|\vec{x}_{\text{opt}} - \mathbb{P}_{j,k'}(\vec{x}_{\text{opt}})\| \leq \tilde{C} \cdot 2^{-j}$. Hence, we now have

$$\|\vec{x} - \mathbb{P}_{j,k'}(\vec{x})\| \leq 2 \|\vec{x} - \vec{x}_{\text{opt}}\| + \tilde{C} \cdot 2^{-j}.$$

The result follows. \square

In this paper we are primarily concerned with achieving approximation results akin to Lemma 2 utilizing *compressive measurements*. This will allow us to extend the successful sparse approximation techniques and results of compressive sensing to the recovery of signals which belong to low dimensional submanifolds of \mathbb{R}^D . In order to accomplish this goal we must first propose and then subsequently analyze both a measurement operator and an associated recovery algorithm. Furthermore, in order for it to be of practical value, we must demonstrate that the proposed recovery algorithm is computationally efficient, easy to implement, and provably accurate. We begin this process by considering our measurement matrices in Section 2.3. We then develop a practical reconstruction algorithm in Section 3. Before we begin, however, we will first state the main results proven herein.

2.2. Main Results. In the statements of the two propositions below, $C \in \mathbb{R}^+$ is an absolute universal constant which is independent of \vec{x} , \mathcal{M} , \mathcal{M} 's GMRA approximation, etcetera. Note that the upper bounds provided for this constant in Section 3 are almost surely quite loose. We state our first result.

Proposition 1. Fix precision parameter $\delta \in \mathbb{R}^+$ and let $\vec{x} \in \mathbb{R}^D$. In addition, let \mathbb{P}_J , $J = O(\log[1/(\delta \text{reach}(\mathcal{M}))])$, be a GMRA approximation to a given compact d -dimensional Riemannian manifold, $\mathcal{M} \subset \mathbb{R}^D$, with volume V . Finally, let

$$m = O \left(d \log \left(\frac{d}{\delta \text{reach}(\mathcal{M})} \right) + \log V \right)$$

be a natural number, and define $\mathcal{A} : \mathbb{R}^m \rightarrow \mathbb{R}^D$ to be Algorithm 1 from Section 3 below. Then, there exists an $m \times D$ matrix, M , such that

$$\|\vec{x} - \mathcal{A}(M\vec{x})\| < C \cdot \|\vec{x} - \vec{x}_{\text{opt}}\| + \delta$$

with arbitrarily high probability. Furthermore, $\mathcal{A}(M\vec{x})$ can be evaluated in $(m^{O(1)} + O(dD))$ -time.

Proof: The result follows from Theorem 3, the first part of Theorem 2, and the discussion in Section 3.1. \square

Proposition 1 provides a nonuniform recovery guarantee for each given $\vec{x} \in \mathbb{R}^D$. If desired, bounds could be altered to depend on the desired probability of success, $p \in (0, 1)$, by including an additional multiplicative factor of $O(\log(1/(1-p)))$ in both the runtime of the algorithm and the upper bound for m . The measurement matrices, M , referred to by the proposition can be any standard Johnson-Lindenstrauss embedding matrix (e.g., a Gaussian random matrix, a random orthogonal projection, etc.). Hence, they are well understood. The worst case theoretical runtime complexity of the recovery algorithm is polynomial in m . We refer the reader to Section 5 for an empirical evaluation of the recovery algorithm's computational efficiency in practice. Finally, we note that the number of required measurements, m , is entirely independent of the extrinsic dimension, D . Next, we state a uniform approximation guarantee for Algorithm 1.

Proposition 2. Fix precision parameter $\delta \in \mathbb{R}^+$. In addition, let \mathbb{P}_J , $J = O(\log[1/(\delta \text{reach}(\mathcal{M}))])$, be a GMRA approximation to a given compact d -dimensional Riemannian manifold, $\mathcal{M} \subset \mathbb{R}^D$, with volume V . Finally, let

$$m = O\left(d \log\left(\frac{D}{\delta \text{reach}(\mathcal{M})}\right) + \log V\right)$$

be a natural number, and define $\mathcal{A} : \mathbb{R}^m \rightarrow \mathbb{R}^D$ to be Algorithm 1 from Section 3 below. Then, there exists an $m \times D$ matrix, M , such that

$$\|\vec{x} - \mathcal{A}(M\vec{x})\| < C \|\vec{x} - \vec{x}_{\text{opt}}\|_2 + \frac{C}{\sqrt{d}} \|\vec{x} - \vec{x}_{\text{opt}}\|_1 + \delta$$

for all $\vec{x} \in \mathbb{R}^D$ with

$$2 \|\vec{x} - \vec{x}_{\text{opt}}\|_2 + \frac{6}{5\sqrt{d}} \|\vec{x} - \vec{x}_{\text{opt}}\|_1 \leq \max\left\{\|\vec{x} - \vec{c}_{J,k_j(\vec{x})}\|, \delta\right\}.$$

Furthermore, $\mathcal{A}(M\vec{x})$ can be evaluated in worst case $(2^{O(d)} \log V + O(md^2 + dD))$ -time.

Proof: The result follows from Theorem 4, the second part of Theorem 2, and the discussion in Section 3.1. \square

Proposition 2 is best interpreted as a general stability result. It guarantees that Algorithm 1 will uniformly approximate all points which are sufficiently close to the manifold \mathcal{M} (i.e., the points need not be exactly on \mathcal{M}). Thus, Algorithm 1 has some limited robustness to arbitrary additive input noise. The examples in the experimental section suggest that the constants involved are very mild.

2.3. The Measurement Matrix. In the process of developing an algorithm to approximate $\mathbb{P}_{j,k_j(\vec{x})}(\vec{x})$, and subsequently demonstrating its accuracy, we will require some knowledge regarding our $m \times D$ measurement matrix M . We shall consider two sets of assumptions regarding M 's interaction with both the manifold \mathcal{M} and our given set of affine projectors for \mathcal{M} at each scale $j \in [J]$. Each set of assumptions will ultimately result in both different approximation guarantees for our reconstruction algorithm, and different measurements bounds (i.e., sufficient upper bounds on m) for M . We will postpone discussion of how to create M and how to bound the number of rows it must have in order to satisfy each set of assumptions below until Section 4. In Section 3 below we will begin by presenting our reconstruction algorithm together with approximation error bounds under each set of assumptions regarding M .

Let $\vec{x} \in \mathbb{R}^D$ and $\mathbb{P} = \{\mathbb{P}_j \mid j \in [J]\}$ be a fixed set of affine projectors for \mathcal{M} for each scale $j \in [J]$. Fix $\epsilon \in (0, \frac{1}{2})$. In Sections 3 and 4 we will assume that our $m \times D$ measurement matrix M satisfies each of these sets of assumptions in turn.

(1) **Assumption Set 1: Required for Nonuniform Recovery of a Given $\vec{x} \in \mathbb{R}^D$ (see Proposition 1)**

(a) Let $S_1 \subset \mathbb{R}^D$ be

$$S_1 = \left\{ \Phi_{j,k}^T \Phi_{j,k} (\vec{x} - \vec{c}_{j,k}) \mid j \in [J], k \in [K_j] \right\} \cup \left\{ \vec{x} - \vec{c}_{j,k} \mid j \in [J], k \in [K_j] \right\} \cup \{\vec{0}\}.$$

We will assume that

$$(1 - \epsilon) \|\vec{y} - \vec{z}\|^2 \leq \|M\vec{y} - M\vec{z}\|^2 \leq (1 + \epsilon) \|\vec{y} - \vec{z}\|^2$$

for all $\vec{y}, \vec{z} \in S_1$.

(b) Furthermore, we will assume that

$$(1 - \epsilon) \|\Phi_{j,k}^T \Phi_{j,k} \vec{y}\| \leq \|M \Phi_{j,k}^T \Phi_{j,k} \vec{y}\| \leq (1 + \epsilon) \|\Phi_{j,k}^T \Phi_{j,k} \vec{y}\|$$

for all $j \in [J]$, $k \in [K_j]$, and $\vec{y} \in \mathbb{R}^D$.

(2) **Assumption Set 2: Required for General Stability (see Proposition 2)**

(a) Let $S_2 = \mathcal{M} \cup \{\vec{c}_{j,k} \mid j \in [J], k \in [K_j]\} \subset \mathbb{R}^D$. We will assume that

$$(1 - \epsilon) \|\vec{y} - \vec{z}\|^2 \leq \|M\vec{y} - M\vec{z}\|^2 \leq (1 + \epsilon) \|\vec{y} - \vec{z}\|^2$$

for all $\vec{y}, \vec{z} \in S_2$.

(b) Furthermore, we will assume that $\|M\vec{y}\|$ is bounded above by $E_M(\vec{y})$ for all $\vec{y} \in \mathbb{R}^D$, where $E_M : \mathbb{R}^D \rightarrow \mathbb{R}^+$ is a continuous function with $E_M(\vec{0}) = 0$. E_M is discussed in detail in Section 4.2.

(c) As before, we will assume that

$$(1 - \epsilon) \|\Phi_{j,k}^T \Phi_{j,k} \vec{y}\| \leq \|M \Phi_{j,k}^T \Phi_{j,k} \vec{y}\| \leq (1 + \epsilon) \|\Phi_{j,k}^T \Phi_{j,k} \vec{y}\|$$

for all $j \in [J]$, $k \in [K_j]$, and $\vec{y} \in \mathbb{R}^D$.

(d) Finally, we will also assume that

$$(1 - \epsilon) \|\vec{y} - \mathbb{P}_{j,k}(\vec{y})\| - 2^{-J} \leq \|M\vec{y} - M\mathbb{P}_{j,k}(\vec{y})\| \leq (1 + \epsilon) \|\vec{y} - \mathbb{P}_{j,k}(\vec{y})\| + 2^{-J}$$

for all $j \in [J]$, $k \in [K_j]$, and $\vec{y} \in \mathcal{M}$.

Note that the critical difference between the two sets of assumptions above concerns the treatment of $\vec{x} \in \mathbb{R}^D$ and $\vec{x}_{\text{opt}} \in \mathcal{M} \subset \mathbb{R}^D$. If possible we would like to obtain measurement bounds which are independent of the ambient dimension, D . Since an arbitrary vector \vec{x} may contain a substantial portion of its energy in the subspace orthogonal to the tangent space to \mathcal{M} at \vec{x}_{opt} , results which are entirely independent of D generally appear to be unattainable unless our measurement matrix happens to successfully preserve information in the direction of $\vec{x} - \vec{x}_{\text{opt}}$. We assume that M preserves lengths of vectors in the general direction of $\vec{x} - \vec{x}_{\text{opt}}$ as part of our first set of assumptions. In the second set of assumptions we do not. It is primarily this difference which leads to different measurement bounds and error guarantees in each case.

3. THE RECONSTRUCTION ALGORITHM

We will ultimately upper bound the number of measurements required in order to approximate a given $\vec{x} \in \mathbb{R}^D$ which is close to $\mathcal{M} \subset \mathbb{R}^D$ via the simple reconstruction technique presented in this section. In doing so we will require that the reconstruction algorithm approximates \vec{x} nearly as well as the vector on \mathcal{M} closest to \vec{x} ,

$$\vec{x}_{\text{opt}} = \arg \min_{\vec{y} \in \mathcal{M}} \|\vec{x} - \vec{y}\|,$$

approximates \vec{x} . Our first order of business, therefore, will be to derive explicit error guarantees for the reconstruction technique considered herein which demonstrate that it is indeed “near-optimal” in the sense discussed in Section 1 above. Let $\mathcal{A}(M\vec{x}) \in \mathbb{R}^D$ denote the output of our reconstruction procedure for a given input $\vec{x} \in \mathbb{R}^D$. We wish to bound the approximation error

$$\|\vec{x} - \mathcal{A}(M\vec{x})\|$$

in terms of the optimal approximation error, $\|\vec{x} - \vec{x}_{\text{opt}}\|$, and an additive error term of size $O(2^{-J})$ whenever possible. Before this task can be accomplished, however, we must first describe the recovery algorithm we will use to calculate $\mathcal{A}(M\vec{x})$.

Our reconstruction procedure uses compressive measurements of \vec{x} in order to approximate $\mathbb{P}_{j,k_j(\vec{x})}(\vec{x})$ in two steps (see Algorithm 1 above). First, the compressive measurements of \vec{x} are used to determine a “center” vector, $\vec{c}_{j,k'}$, which is nearly as close to \vec{x} as its nearest neighboring center, $\vec{c}_{j,k_j(\vec{x})}$, is. This step is guaranteed to work well as long as our measurement matrix, M , preserves appropriate distances between \vec{x} and all the center vectors at scale j .

Algorithm 1 APPROXIMATE $\mathbb{P}_{j,k_j(\vec{x})}(\vec{x})$

- 1: **Input:** Measurements $M\vec{x} \in \mathbb{R}^m$, **Scale** $j \in [J]$, **Approximation** $\mathbb{P}_j = \{\mathbb{P}_{j,k} \mid k \in [K_j]\}$ **to manifold** $\mathcal{M} \subset \mathbb{R}^D$
 - 2: **Output:** $\mathcal{A}(M\vec{x})$, an approximation to $\mathbb{P}_{j,k_j(\vec{x})}(\vec{x}) \approx \vec{x}$
 - 3: $k' \leftarrow \arg \min_{k \in [K_j]} \|M\vec{x} - M\vec{c}_{j,k}\|$
 - 4: $\vec{u}' \leftarrow \arg \min_{\vec{u} \in \mathbb{R}^d} \|M\Phi_{j,k'}^T \vec{u} - M\vec{x} + M\vec{c}_{j,k'}\|$
 - 5: $\mathcal{A}(M\vec{x}) \leftarrow \Phi_{j,k'}^T \vec{u}' + \vec{c}_{j,k'}$
 - 6: Output $\mathcal{A}(M\vec{x})$
-

Next, an accurate projection of $\vec{x} - \vec{c}_{j,k'}$ onto the d -dimensional subspace associated with $\vec{c}_{j,k'}$ is found by solving an overdetermined least squares problem. This step will also work well as long as our measurement matrix M is well conditioned on all of the d -dimensional subspaces associated with the scale j center vectors. As we demonstrate below, the two sets of assumptions for M in Section 2.3 are sufficient to guarantee that both steps work well.

The following lemma guarantees that the center found in line 3 of Algorithm 1 is nearly as close to \vec{x} as \vec{x} 's true nearest center is.

Lemma 3. Fix $\epsilon \in (0, \frac{1}{2})$. Let $\mathcal{M} \subset \mathbb{R}^D$ be a compact d -dimensional Riemannian submanifold of \mathbb{R}^D , and $\vec{x} \in \mathbb{R}^D$. Furthermore, let $\mathbb{P}_j = \{\mathbb{P}_{j,k} \mid k \in [K_j]\}$ be a scale $j \in [J]$ GWRA approximation to \mathcal{M} . Then, if our $m \times D$ measurement matrix M satisfies Assumption Set 1 in Section 2.3 above, line 3 of Algorithm 1 will select a $k' \in [K_j]$ which has

$$\|\vec{x} - \vec{c}_{j,k'}\| \leq \sqrt{\frac{1+\epsilon}{1-\epsilon}} \cdot \|\vec{x} - \vec{c}_{j,k_j(\vec{x})}\|.$$

If our $m \times D$ measurement matrix M satisfies Assumption Set 2 in Section 2.3 above, then line 3 of Algorithm 1 will select a $k' \in [K_j]$ which has

$$(4) \quad \|\vec{x} - \vec{c}_{j,k'}\| \leq \sqrt{\frac{1+\epsilon}{1-\epsilon}} \cdot \|\vec{x} - \vec{c}_{j,k_j(\vec{x})}\| + \left(1 + \sqrt{\frac{1+\epsilon}{1-\epsilon}}\right) \|\vec{x} - \vec{x}_{\text{opt}}\| + \sqrt{\frac{4}{1-\epsilon}} \cdot E_M(\vec{x} - \vec{x}_{\text{opt}}).$$

Proof: Using the first set of assumptions for M together with the definition of $k' \in [K_j]$ from Algorithm 1 we can see that

$$\|\vec{x} - \vec{c}_{j,k'}\| \leq \sqrt{\frac{1}{1-\epsilon}} \cdot \|M\vec{x} - M\vec{c}_{j,k'}\| \leq \sqrt{\frac{1}{1-\epsilon}} \cdot \|M\vec{x} - M\vec{c}_{j,k_j(\vec{x})}\| \leq \sqrt{\frac{1+\epsilon}{1-\epsilon}} \cdot \|\vec{x} - \vec{c}_{j,k_j(\vec{x})}\|.$$

We now turn our attention to the case where M satisfies the second set of assumptions. We have that

$$\begin{aligned} \|\vec{x} - \vec{c}_{j,k'}\| &\leq \|\vec{x} - \vec{x}_{\text{opt}}\| + \|\vec{x}_{\text{opt}} - \vec{c}_{j,k'}\| \leq \|\vec{x} - \vec{x}_{\text{opt}}\| + \sqrt{\frac{1}{1-\epsilon}} \cdot \|M\vec{x}_{\text{opt}} - M\vec{c}_{j,k'}\| \\ &\leq \|\vec{x} - \vec{x}_{\text{opt}}\| + \sqrt{\frac{1}{1-\epsilon}} \left(\|M(\vec{x} - \vec{x}_{\text{opt}})\| + \|M\vec{x} - M\vec{c}_{j,k_j(\vec{x})}\| \right) \\ &\leq \|\vec{x} - \vec{x}_{\text{opt}}\| + \sqrt{\frac{4}{1-\epsilon}} \cdot \|M(\vec{x} - \vec{x}_{\text{opt}})\| + \sqrt{\frac{1+\epsilon}{1-\epsilon}} \cdot \|\vec{x}_{\text{opt}} - \vec{c}_{j,k_j(\vec{x})}\|. \end{aligned}$$

Focusing on the first and third terms in the line immediately above, we note that

$$\|\vec{x} - \vec{x}_{\text{opt}}\| + \sqrt{\frac{1+\epsilon}{1-\epsilon}} \cdot \|\vec{x}_{\text{opt}} - \vec{c}_{j,k_j(\vec{x})}\| \leq \sqrt{\frac{1+\epsilon}{1-\epsilon}} \cdot \|\vec{x} - \vec{c}_{j,k_j(\vec{x})}\| + \left(1 + \sqrt{\frac{1+\epsilon}{1-\epsilon}}\right) \|\vec{x} - \vec{x}_{\text{opt}}\|.$$

The result follows. \square

Next, we prove a lemma which guarantees the accuracy of the solution of the overdetermined least squares problem produced by line 4 of Algorithm 1.

Lemma 4. Let $\mathcal{M} \subset \mathbb{R}^D$ be a compact d -dimensional Riemannian submanifold of \mathbb{R}^D , and $\vec{x} \in \mathbb{R}^D$. Furthermore, let $\mathbb{P}_j = \{\mathbb{P}_{jk} \mid k \in [K_j]\}$ be a scale $j \in [J]$ GWRA approximation to \mathcal{M} , and $k' \in [K_j]$ be the value computed by line 3 of Algorithm 1. Then, if our $m \times D$ measurement matrix M satisfies either set of assumptions in Section 2.3 above, line 5 of Algorithm 1 will produce an $\mathcal{A}(M\vec{x}) \in \mathbb{R}^D$ which has

$$\|\mathbb{P}_{jk'}(\vec{x}) - \mathcal{A}(M\vec{x})\| \leq \frac{2}{1-\epsilon} \cdot \|M[\vec{x} - \mathbb{P}_{jk'}(\vec{x})]\|.$$

Proof: Let $\vec{u}' \in \mathbb{R}^D$ be as defined in line 4 of Algorithm 1. Given either set of assumptions for M we will have

$$\|\Phi_{jk'}^T \vec{u}' - \Phi_{jk'}^T \Phi_{jk'}(\vec{x} - \vec{c}_{jk'})\| \leq \frac{1}{1-\epsilon} \left(\|M\Phi_{jk'}^T \vec{u}' - M(\vec{x} - \vec{c}_{jk'})\| + \|M[I - \Phi_{jk'}^T \Phi_{jk'}](\vec{x} - \vec{c}_{jk'})\| \right),$$

where I is the $D \times D$ identity matrix. By the definition of \vec{u}' in Algorithm 1 we can now see that

$$\|\Phi_{jk'}^T \vec{u}' - \Phi_{jk'}^T \Phi_{jk'}(\vec{x} - \vec{c}_{jk'})\| \leq \frac{2}{1-\epsilon} \cdot \|M[I - \Phi_{jk'}^T \Phi_{jk'}](\vec{x} - \vec{c}_{jk'})\|.$$

The stated result follows. \square

Finally, we demonstrate the accuracy of the output of Algorithm 1 as an approximation to \vec{x} .

Theorem 2. Fix $\epsilon \in (0, \frac{1}{2})$. Let $\mathcal{M} \subset \mathbb{R}^D$ be a compact d -dimensional Riemannian submanifold of \mathbb{R}^D , and $\vec{x} \in \mathbb{R}^D$. Furthermore, let $\mathbb{P}_j = \{\mathbb{P}_{jk} \mid k \in [K_j]\}$ be a scale $j \in [J]$ GWRA approximation to \mathcal{M} . Then, if our $m \times D$ measurement matrix M satisfies Assumption Set 1 in Section 2.3 above, Algorithm 1 will output a point, $\mathcal{A}(M\vec{x}) \in \mathbb{R}^D$, which satisfies

$$\|\vec{x} - \mathcal{A}(M\vec{x})\| < 100.3 \|\vec{x} - \vec{x}_{\text{opt}}\| + O(2^{-j}).$$

Now suppose that our $m \times D$ measurement matrix M satisfies Assumption Set 2 in Section 2.3 above, and that \mathbb{P}_j is a scale j GWRA approximation to \mathcal{M} for some $j > j_0$ (revisit Properties 3a and 3b in Section 2 for the definitions of the constants j_0 , C_1 , and \tilde{C}). Furthermore, suppose that $\vec{x} \in \mathbb{R}^D - \mathcal{M}$ has

$$(5) \quad 2 \cdot E_M(\vec{x} - \vec{x}_{\text{opt}}) \leq (8\sqrt{1-\epsilon} - \sqrt{1+\epsilon}) \|\vec{x} - \vec{c}_{jk_j(\vec{x})}\| - (\sqrt{1-\epsilon} + \sqrt{1+\epsilon}) \|\vec{x} - \vec{x}_{\text{opt}}\|.$$

Then, Algorithm 1 will output a point, $\mathcal{A}(M\vec{x}) \in \mathbb{R}^D$, which satisfies

$$(6) \quad \|\vec{x} - \mathcal{A}(M\vec{x})\| < 220 \|\vec{x} - \vec{x}_{\text{opt}}\| + 4 \cdot E_M(\vec{x} - \vec{x}_{\text{opt}}) + O(2^{-j}).$$

Proof: To begin we note that

$$(7) \quad \|\vec{x} - \mathcal{A}(M\vec{x})\| \leq \|\vec{x} - \mathbb{P}_{jk'}(\vec{x})\| + \|\mathbb{P}_{jk'}(\vec{x}) - \mathcal{A}(M\vec{x})\|$$

where $k' \in [K_j]$ is defined as in line 3 of Algorithm 1. The first set of assumptions for M together with Lemmas 2 and 3 tells us that

$$\|\vec{x} - \mathbb{P}_{jk'}(\vec{x})\| \leq 17 \|\vec{x} - \vec{x}_{\text{opt}}\| + O(2^{-j})$$

since $\epsilon \in (0, \frac{1}{2})$. Furthermore, the first set of assumptions for M together with Lemma 4 indicates that

$$\|\mathbb{P}_{jk'}(\vec{x}) - \mathcal{A}(M\vec{x})\| \leq \frac{2}{1-\epsilon} \cdot \|M[\vec{x} - \mathbb{P}_{jk'}(\vec{x})]\| \leq 2 \frac{\sqrt{1+\epsilon}}{1-\epsilon} \cdot \|\vec{x} - \mathbb{P}_{jk'}(\vec{x})\|.$$

Hence, we obtain the stated bound in the first case.

Now assume that M satisfies Assumption Set 2 in Section 2.3. We will begin by bounding the $\|\mathbb{P}_{jk'}(\vec{x}) - \mathcal{A}(M\vec{x})\|$ term in Equation 7. Applying Lemma 4 and then utilizing our second set of assumptions regarding M we can see that

$$(8) \quad \begin{aligned} \|\mathbb{P}_{jk'}(\vec{x}) - \mathcal{A}(M\vec{x})\| &\leq \frac{2}{1-\epsilon} \cdot \|M[\vec{x} - \mathbb{P}_{jk'}(\vec{x})]\| \leq \frac{2}{1-\epsilon} (\|M\vec{x} - M\vec{x}_{\text{opt}}\| + \|M\vec{x}_{\text{opt}} - M\mathbb{P}_{jk'}(\vec{x})\|) \\ &\leq \frac{2}{1-\epsilon} (\|M\vec{x} - M\vec{x}_{\text{opt}}\| + \|M\vec{x}_{\text{opt}} - M\mathbb{P}_{jk'}(\vec{x}_{\text{opt}})\| + \|M\mathbb{P}_{jk'}(\vec{x}_{\text{opt}}) - M\mathbb{P}_{jk'}(\vec{x})\|) \\ &\leq 2 \cdot \frac{1+\epsilon}{1-\epsilon} \cdot \|\vec{x} - \vec{x}_{\text{opt}}\| + \frac{2}{1-\epsilon} (\|M\vec{x} - M\vec{x}_{\text{opt}}\| + \|M\vec{x}_{\text{opt}} - M\mathbb{P}_{jk'}(\vec{x}_{\text{opt}})\|). \end{aligned}$$

In order to bound the last term in Equation 8 above, we note that $\|\vec{x} - \vec{c}_{j,k'}\| \leq 8 \|\vec{x} - \vec{c}_{j,k_j(\vec{x})}\|$ whenever $E_M(\vec{x} - \vec{x}_{\text{opt}})$ satisfies Equation 5. Therefore, we will have $\|\vec{x}_{\text{opt}} - \vec{c}_{j,k'}\| < 16 \|\vec{x}_{\text{opt}} - \vec{c}_{j,k_j(\vec{x}_{\text{opt}})}\|$ whenever $\|\vec{x} - \vec{c}_{j,k'}\| > 17 \|\vec{x} - \vec{x}_{\text{opt}}\|$ by an argument identical to that presented in the second paragraph of the proof of Lemma 2. Hence, Property 3b in Section 2 guarantees that $\|\vec{x}_{\text{opt}} - \mathbb{P}_{j,k'}(\vec{x}_{\text{opt}})\| \leq \tilde{C} \cdot 2^{-j}$ whenever $\|\vec{x} - \vec{c}_{j,k'}\| > 17 \|\vec{x} - \vec{x}_{\text{opt}}\|$. Item (d) of Assumption Set 2 in Section 2.3 now guarantees that $\|M\vec{x}_{\text{opt}} - M\mathbb{P}_{j,k'}(\vec{x}_{\text{opt}})\|$ will also be $O(2^{-j})$ whenever $\|\vec{x} - \vec{c}_{j,k'}\| > 17 \|\vec{x} - \vec{x}_{\text{opt}}\|$.

To finish, suppose that $\|\vec{x} - \vec{c}_{j,k'}\| \leq 17 \|\vec{x} - \vec{x}_{\text{opt}}\|$. Continuing to bound the last term of Equation 8 in this case we obtain

$$\begin{aligned} \frac{2}{1-\epsilon} \|M\vec{x}_{\text{opt}} - M\mathbb{P}_{j,k'}(\vec{x}_{\text{opt}})\| &\leq \frac{2}{1-\epsilon} \left(\|M\vec{x}_{\text{opt}} - M\vec{c}_{j,k'}\| + \|M\Phi_{j,k'}^T \Phi_{j,k'}(\vec{x}_{\text{opt}} - \vec{c}_{j,k'})\| \right) \\ &\leq 2 \cdot \frac{\sqrt{1+\epsilon}}{1-\epsilon} \|\vec{x}_{\text{opt}} - \vec{c}_{j,k'}\| + 2 \cdot \frac{1+\epsilon}{1-\epsilon} \|\vec{x}_{\text{opt}} - \vec{c}_{j,k'}\| \\ &\leq 36 \cdot \frac{\sqrt{1+\epsilon}}{1-\epsilon} (1 + \sqrt{1+\epsilon}) \|\vec{x} - \vec{x}_{\text{opt}}\|. \end{aligned}$$

Combining this bound with the previous paragraph concludes the proof. \square

Theorem 2 demonstrates that Algorithm 1 can stably approximate vectors $\vec{x} \in \mathbb{R}^D - \mathcal{M}$ as long as the measurement matrix, M , satisfies one of the two sets of assumptions detailed in Section 2.3. However, the strength of the approximation guarantee depends on which set of assumptions M satisfies. When M possess the attributes listed in Assumption Set 1 (most notably, attribute (a)) the vector returned by Algorithm 1 will always provide an approximation to \vec{x} whose error is within a constant multiple of the optimal approximation error. When M satisfies Assumption Set 2, on the other hand, Algorithm 1 is only guaranteed to provide near optimal approximations for vectors, \vec{x} , which are relatively close to the manifold \mathcal{M} .

3.1. Practical Implementation of Algorithm 1. In line 3 of Algorithm 1 we want to locate the nearest neighbor of $M\vec{x} \in \mathbb{R}^m$ from the set $\{M\vec{c}_{j,k} \mid k \in [K_j]\} \subset \mathbb{R}^m$. This can be accomplished naively in $O(mK_j)$ -time. However, K_j is potentially large in the worst case (see Lemma 6 below). Therefore, it is important to note that the runtime's dependence on K_j can be greatly reduced in practice with the aid of standard space partitioning techniques (e.g., by building a k-d tree to solve the nearest neighbor problem). Alternatively, other fast nearest neighbor methods could also be utilized (e.g., see [31, 8, 4] and the references therein). Due to the dyadic structure of our $\vec{c}_{j,k}$ -vectors, the worst case theoretical runtime complexity of line 3 can be improved slightly to $(2^{O(d)} \log V)$ -time by using cover trees [8].² Alternatively, if it suffices to find a $(1 + \delta)$ -nearest neighbor of $M\vec{x}$ with high probability, we can utilize even faster algorithms which run in $m^{O(1)}$ -time (see Proposition 3 in [31] together with the bound for m in Theorem 3 below).

Line 4 of Algorithm 1 requires the solution of an overdetermined least squares problem. This can be accomplished in $O(md^2)$ -time via the singular value decomposition of $M\Phi_{j,k'}^T$. Furthermore, the solution can be computed accurately since both sets of assumptions in Section 2.3 guarantee that $M\Phi_{j,k'}^T$ is well conditioned. Finally, explicitly forming $\mathcal{A}(M\vec{x})$ in line 5 of Algorithm 1 can be accomplished in $O(Dd)$ -time. The total runtime of Algorithm 1 will therefore be $O(d(md + D) + T_{\text{NN}})$, where T_{NN} bounds the runtime of the nearest neighbor algorithm used in line 3.

4. UPPER BOUNDS ON THE NUMBER OF REQUIRED MEASUREMENTS

In this section we will bound the number of rows, m , needed in order for our $m \times N$ measurement matrix, M , to satisfy each set of assumptions discussed in Section 2.3. In order to do so, it will suffice to let M be a linear Johnson-Lindenstrauss embedding of a well chosen set of points in \mathbb{R}^D into \mathbb{R}^m . Of course, this set of points will vary depending on which set of assumptions from Section 2.3 we want M to satisfy. Below we consider each set of assumptions separately. However, we will first establish two lemmas which will be useful in both cases.

Lemma 5. *Let $\epsilon \in (0, \frac{1}{2})$. Furthermore, let $j \in [J]$ and $k \in [K_j]$ denote an affine projector $\mathbb{P}_{j,k}$ (see Property 1 in Section 2). Then, there exists a finite set of vectors, $Q_{j,k} \subset X_{j,k} = \{\Phi_{j,k}^T \Phi_{j,k} \vec{y} \mid \vec{y} \in \mathbb{R}^D\}$ with $|Q_{j,k}| \leq (12/\epsilon)^d + 1$, such*

²Here V is the volume of the d -dimensional manifold $\mathcal{M} \subset \mathbb{R}^D$.

that

$$(1 - \epsilon) \left\| \Phi_{j,k}^T \Phi_{j,k} \vec{y} \right\| \leq \left\| M \Phi_{j,k}^T \Phi_{j,k} \vec{y} \right\| \leq (1 + \epsilon) \left\| \Phi_{j,k}^T \Phi_{j,k} \vec{y} \right\|$$

for all $\vec{y} \in \mathbb{R}^D$ whenever M embeds $Q_{j,k}$ into \mathbb{R}^m with $\epsilon/2$ -distortion.

Proof: We let $Q'_{j,k}$ be a minimal $\epsilon/4$ -cover of the d -dimensional unit ball in $X_{j,k}$ centered at $\vec{0} \in X_{j,k}$. Now set $Q_{j,k} = Q'_{j,k} \cup \{\vec{0}\}$. The stated upper bound of $|Q_{j,k}|$ follows from existing covering results (see [5] for references). Furthermore, if M embeds $Q_{j,k}$ into \mathbb{R}^m with $\epsilon/2$ -distortion it is easy to see that

$$(1 - \epsilon/2) \|q\| \leq \|Mq\| \leq (1 + \epsilon/2) \|q\|$$

for all $q \in Q_{j,k}$. The remainder of the proof now directly parallels the proof of Lemma 5.1 in [5]. \square

Lemma 6. Fix $J \in \mathbb{N}$ and let \mathbb{P}_j , $j \in [J]$, be a GMRA approximation to a given compact d -dimensional Riemannian manifold, $\mathcal{M} \subset \mathbb{R}^D$, with d -dimensional volume V . Furthermore, suppose that $j' \in [J] - \left[\max \left\{ j_0, \log_2 \left(\frac{C_1}{\text{reach}(\mathcal{M})} \right) - 2 \right\} \right]$, where j_0 and C_1 are defined as in Property 3a of Section 2. Then, the number of affine projectors at scale j' , $K_{j'}$, is bounded above by $\frac{2^{d(j'+1.5)}}{C_1^d} \cdot V \cdot \left(\frac{d}{2} + 1 \right)^{\frac{d}{2}+1}$.

Proof: We know that $\mathcal{B}_{C_1 \cdot 2^{-j'-2}}(\vec{c}_{j',k}) \cap \mathcal{M}$ is nonempty for all $k \in [K_{j'}]$ since $j' > j_0$. Now consider a minimal $C_1 \cdot 2^{-j'-2}$ -cover of \mathcal{M} , $C_{C_1 \cdot 2^{-j'-2}}(\mathcal{M})$. It is not difficult to see that every $\vec{c}_{j',k}$ will be contained in $\mathcal{B}_{C_1 \cdot 2^{-j'-1}}(\vec{y})$ for some $\vec{y} \in C_{C_1 \cdot 2^{-j'-2}}(\mathcal{M})$. Furthermore, there can be no $\vec{y} \in C_{C_1 \cdot 2^{-j'-2}}(\mathcal{M})$ such that two distinct $\vec{c}_{j',k}$ are contained in the same ball, $\mathcal{B}_{C_1 \cdot 2^{-j'-1}}(\vec{y})$, by Property 2b in Section 2. Hence, $K_{j'} \leq |C_{C_1 \cdot 2^{-j'-2}}(\mathcal{M})|$. Applying Lemma 1 concludes the proof. \square

We are now prepared to upper bound the number of rows required by our $m \times N$ measurement matrix, M , in order to satisfy each set of assumptions listed in Section 2.3.

4.1. Bounding the Number of Rows Required to Satisfy Assumption Set 1.

Theorem 3. Fix $\epsilon \in (0, \frac{1}{2})$, $\vec{x} \in \mathbb{R}^D$, and $J \in \mathbb{N}$ sufficiently large. Furthermore, let \mathbb{P}_j , $j \in [J]$, be a GMRA approximation to a given compact d -dimensional Riemannian manifold, $\mathcal{M} \subset \mathbb{R}^D$, with volume V . Then, there exists an $m \times D$ matrix, M , which satisfies Assumption Set 1 in Section 2.3 with $m = O(d\epsilon^{-2}(J + \log(d/\epsilon)) + \epsilon^{-2} \log V)$.

Proof: The set $S_1 \subset \mathbb{R}^D$ defined in item (a) of Assumption Set 1 has $|S_1| \leq 2(J + 1)K_J + 1$. Furthermore, applying Lemma 5 to all at most $(J + 1)K_J$ affine projectors yields a set of size at most $(J + 1)K_J \left((12/\epsilon)^d + 1 \right)$ for item (b) of Assumption Set 1. Lemma 6 together with Theorem 1 now finishes the proof. \square

It is important to recall that Theorem 1 is proven by showing that a random matrix will (nearly) isometrically embed a given subset of \mathbb{R}^D into \mathbb{R}^m with high probability. In the proof of Theorem 3 above, Theorem 1 is applied to embed a set which depends on the given $\vec{x} \in \mathbb{R}^D$ we are ultimately interested in approximating (i.e., the set S_1 defined in Section 2.3 depends on \vec{x}). Thus, Theorem 3 provides us with a high probability recovery guarantee for each separate $\vec{x} \in \mathbb{R}^D$ on which we apply Algorithm 1.

4.2. Bounding the Number of Rows Required to Satisfy Assumption Set 2. We will begin this section by considering item (b) of Assumption Set 2. Among other things, this will allow us to finally define the function $E_M : \mathbb{R}^D \rightarrow \mathbb{R}^+$. However, we must first define the Restricted Isometry Property [12] on which the subsequent discussion relies.

Definition 2. Let $D, d \in \mathbb{N}$, and $\epsilon \in (0, 1)$. An $m \times D$ matrix M' has the Restricted Isometry Property, $RIP(D, d, \epsilon)$, if

$$(9) \quad (1 - \epsilon) \|\vec{x}\|^2 \leq \|M' \vec{x}\|^2 \leq (1 + \epsilon) \|\vec{x}\|^2$$

for all $\vec{x} \in \mathbb{R}^D$ containing at most d nonzero coordinates.

We have the following lemma.

Lemma 7. Let $\epsilon \in (0, \frac{1}{2})$. There exists a finite set of vectors, $Q \subset X = \{\vec{y} \mid \vec{y} \in \mathbb{R}^D \text{ contains } d \text{ nonzero coordinates}\}$ with $|Q| \leq \binom{D}{d} ((12/\epsilon)^d + 1)$, such that an $m \times D$ matrix M' has the $\text{RIP}(D, d, \epsilon)$ whenever it embeds Q into \mathbb{R}^m with $\epsilon/2$ -distortion. Furthermore, any such matrix M' will have $\|M' \vec{y}\|_2$ bounded above by

$$E_{M'}(\vec{y}) = \sqrt{1 + \epsilon} \cdot \left[\|\vec{y}\|_2 + \frac{1}{\sqrt{d}} \|\vec{y}\|_1 \right]$$

for all $\vec{y} \in \mathbb{R}^D$.

Proof: To prove that M' has the $\text{RIP}(D, d, \epsilon)$ we employ an argument similar to the proof of Theorem 5.2 in [5]. To begin, we define \vec{e}_j , $j \in [D] - \{0\}$, to be the j^{th} row of the $D \times D$ identity matrix. Then, for each d -element subset $S = \{j_1, \dots, j_d\} \subset [D] - \{0\}$, we define X_S to be the d -dimensional subspace spanned by $\vec{e}_{j_1}, \dots, \vec{e}_{j_d}$. Next, we let Q'_S be a minimal $\epsilon/4$ -cover of the d -dimensional unit ball in X_S centered at $\vec{0}$, and define $Q_S = Q'_S \cup \{\vec{0}\}$ as per Lemma 5. Finally, we let

$$Q := \bigcup_{S \subset [D] - \{0\}, |S|=d} Q_S.$$

The upper bound on $|Q|$ follows immediately.

Now suppose that M' embeds Q into \mathbb{R}^m with $\epsilon/2$ -distortion. Every $\vec{x} \in \mathbb{R}^D$ containing at most d nonzero coordinates belongs to some subspace, X_S , whose associated set, $Q_S \subset Q$, is also embedded into \mathbb{R}^m with $\epsilon/2$ -distortion by M' . Hence, a trivial variant of Lemma 5 guarantees that every such \vec{x} will satisfy Equation 9. Therefore, M' will have the $\text{RIP}(D, d, \epsilon)$ as claimed. The equation for $E_{M'}$ now follows from Proposition 3.5 in [36]. \square

We are now sufficiently equipped to consider item (a) of Assumption Set 2 in Section 2.3. We have the following lemma.

Lemma 8. Fix $\epsilon \in (0, \frac{1}{2})$ and $J \in \mathbb{N} - [\max\{j_0, \log_2(\frac{C_1}{\text{reach}(\mathcal{M})}) - 2\}]$, where j_0 and C_1 are defined as in Property 3a of Section 2. In addition, let \mathbb{P}_j , $j \in [J]$, be a GMRA approximation to a given compact d -dimensional Riemannian manifold, $\mathcal{M} \subset \mathbb{R}^D$, with d -dimensional volume V . Then, there exist absolute universal constants, $C_3, C_4 \in \mathbb{R}^+$, which are independent of both \mathcal{M} and its GMRA approximation, together with a finite set of vectors, $\tilde{B} \subset \mathbb{R}^D$, so that any $m \times D$ matrix M' which embeds \tilde{B} into \mathbb{R}^m with $(C_3 \cdot \epsilon)$ -distortion will satisfy

$$(1 - \epsilon) \|\vec{y} - \vec{z}\|^2 \leq \|M' \vec{y} - M' \vec{z}\|^2 \leq (1 + \epsilon) \|\vec{y} - \vec{z}\|^2$$

for all $\vec{y}, \vec{z} \in \mathcal{M} \cup \{\vec{c}_{j,k} \mid j \in [J], k \in [K_j]\} \subset \mathbb{R}^D$. Furthermore, $\tilde{B} \subset \mathbb{R}^D$ will have

$$|\tilde{B}| = O\left(2^{C_4 J \cdot d} V^2 \left(\frac{D}{\epsilon \cdot \min\{1, \text{reach}(\mathcal{M})\} \cdot \min\{1, C_1\}}\right)^{C_4 d}\right).$$

Proof: See Appendix A. \square

Furthermore, a modification of the proof of Lemma 8 yields our final lemma concerning Assumption Set 2 in Section 2.3. We have the following result regarding item (d) of Assumption Set 2.

Lemma 9. Fix $\epsilon \in (0, \frac{1}{2})$ and $J \in \mathbb{N} - [\max\{j_0, \log_2(\frac{C_1}{\text{reach}(\mathcal{M})}) - 2\}]$, where j_0 and C_1 are defined as in Property 3a of Section 2. In addition, let \mathbb{P}_j , $j \in [J]$, be a GMRA approximation to a given compact d -dimensional Riemannian manifold, $\mathcal{M} \subset \mathbb{R}^D$, with d -dimensional volume V . Then, there exist absolute universal constants, $C_5, C_6 \in \mathbb{R}^+$, which are independent of both \mathcal{M} and its GMRA approximation, together with a finite set of vectors, $B' \subset \mathbb{R}^D$, so that any $m \times D$ matrix M' which embeds B' into \mathbb{R}^m with $(C_5 \cdot \epsilon)$ -distortion will satisfy

$$(1 - \epsilon) \|\vec{y} - \mathbb{P}_{j,k}(\vec{y})\| - 2^{-J} \leq \|M' \vec{y} - M' \mathbb{P}_{j,k}(\vec{y})\| \leq (1 + \epsilon) \|\vec{y} - \mathbb{P}_{j,k}(\vec{y})\| + 2^{-J}$$

for all $j \in [J]$, $k \in [K_j]$, and $\vec{y} \in \mathcal{M}$. Furthermore, $B' \subset \mathbb{R}^D$ will have

$$|B'| = O\left(2^{C_6 J \cdot d} V^2 \left(\frac{D}{\epsilon \cdot \min\{1, \text{reach}(\mathcal{M})\} \cdot \min\{1, C_1\}}\right)^{C_6 d}\right).$$

Proof: See Appendix B. \square

We are finally ready to provide a useful upper bound for the number of rows required in any measurement matrix satisfying Assumption Set 2 in Section 2.3. We have the following theorem.

Theorem 4. Fix $\epsilon \in (0, \frac{1}{2})$ and $J \in \mathbb{N} - \left[\max \left\{ j_0, \log_2 \left(\frac{C_1}{\text{reach}(\mathcal{M})} \right) - 2 \right\} \right]$, where j_0 and C_1 are defined as in Property 3a of Section 2. In addition, let \mathbb{P}_j , $j \in [J]$, be a GMRA approximation to a given compact d -dimensional Riemannian manifold, $\mathcal{M} \subset \mathbb{R}^D$, with d -dimensional volume V . Then, there exists an $m \times D$ matrix, M , which satisfies Assumption Set 2 in Section 2.3 with

$$m = O \left(d\epsilon^{-2} \log \left(\frac{D}{\epsilon \cdot \text{reach}(\mathcal{M})} \right) + d\epsilon^{-2}J + \epsilon^{-2} \log V \right)$$

and

$$E_M(\vec{y}) = \sqrt{1 + \epsilon} \cdot \left[\|\vec{y}\|_2 + \frac{1}{\sqrt{d}} \|\vec{y}\|_1 \right].$$

Proof: Any $m \times D$ matrix which embeds $\tilde{B} \subset \mathbb{R}^D$ from Lemma 8 into \mathbb{R}^m with $(C_3 \cdot \epsilon)$ -distortion will satisfy both items (a) and (b) of Assumption Set 2 in Section 2.3 (see Lemmas 7 and 8). Similarly, any given $m \times D$ matrix which embeds $B' \subset \mathbb{R}^D$ from Lemma 9 into \mathbb{R}^m with $(C_5 \cdot \epsilon)$ -distortion will satisfy item (d) of Assumption Set 2. Finally, just as in the proof of Theorem 3 above, Lemma 5 applied to all at most $(J + 1)K_J$ affine projectors yields a subset of \mathbb{R}^D of size at most $(J + 1)K_J \left((12/\epsilon)^d + 1 \right)$ for item (c) of Assumption Set 2. Theorem 1 applied to the union of this subset with $\tilde{B} \cup B'$ guarantees the existence of

$$O \left(\epsilon^{-2} \log \left(|\tilde{B}| + |B'| + (J + 1)K_J \left((12/\epsilon)^d + 1 \right) \right) \right) \times D$$

Johnson-Lindenstrauss embedding matrices which satisfy Assumption Set 2 with high probability. Applying Lemmas 6, 8, and 9 to bound K_J , $|\tilde{B}|$, and $|B'|$, respectively, now finishes the proof. \square

In the proof of Theorem 4 above, Theorem 1 is applied to embed a set which only depends on the given manifold, \mathcal{M} , and its GMRA approximation. More specifically, no knowledge was assumed regarding any point $\vec{x} \in \mathbb{R}^D - \mathcal{M}$ which we might be interested in approximating via Algorithm 1. Thus, Theorem 4 provides us with a uniform approximation guarantee for all $\vec{x} \in \mathbb{R}^D$ on which we might apply Algorithm 1. However, we pay several penalties for this uniformity. First, the number of rows in our measurement matrix, m , now depends on the *extrinsic dimensionality*, D , of the given manifold. Second, the resulting uniform error bounds are only nontrivial for input points, \vec{x} , which are close to the given manifold. Hence, although Theorem 4 implies that Algorithm 1 enjoys a limited form of stability, it does not provide very robust uniform error guarantees in practice.

5. EMPIRICAL EVALUATION

We implemented Algorithm 1 and present an empirical evaluation of the algorithm in this section.³ We consider the following examples:

- (i) \mathcal{M}_1 : 20,000 points sampled from a “swiss roll”, a 2-dimensional manifold \mathcal{S} ;
- (ii) \mathcal{M}_2 : 40,000 points sampled from a unit 9-dimensional sphere \mathcal{S}^9 ;
- (iii) \mathcal{M}_3 : 5,000 pictures of the digit ‘1’ from the MNIST data base of images, 28×28 pixels, of handwritten digits⁴, with each picture having pixel intensity normalized to have unit L^2 norm.
- (iii) \mathcal{M}_4 : 15,000 points from the MNIST data base, with 5,000 points sampled from each of the digits 1, 3, 5, with each picture having pixel intensity normalized to have unit L^2 norm.
- (iv) \mathcal{M}_5 : the Science News text document data set, which comprises 1163 text documents, modeled as vectors in 1153 dimensions, whose i -th entry is the frequency of the i -th word in a dictionary (see [19] for detailed information about this data set), normalized so that every document vector has unit Euclidean norm.

³All code is freely available at <http://www.math.duke.edu/~mauro>

⁴Available at <http://yann.lecun.com/exdb/mnist/>.

We construct the GMRA on these data sets in order to obtain the linear approximations, \mathbb{P}_j for each scale j considered below in the noiseless setting.

For the noisy experiments we add Gaussian noise, $\mathcal{N}(0, \frac{\sigma^2}{D} I_D)$ where D is the (ambient) dimension of the data, to each data point for $\sigma = 0, 0.05, 0.1$. We then use the noisy data to compute the GMRA approximations of the noisy data, as well as the random projections utilized by the proposed reconstruction algorithm \mathcal{A} . We consider the following measures of approximation:

$$(10) \quad \text{relMSE}(\mathcal{A}, \mathcal{M}, j)^2 := \frac{1}{n} \sum_{i=1}^n \frac{\|\vec{x}_i - \mathcal{A}(M\vec{x}_i)\|^2}{\|\vec{x}_i\|^2}, \quad \text{relMSE}_j^2 := \frac{1}{n} \sum_{i=1}^n \frac{\|\vec{x}_i - \mathbb{P}_j(\vec{x}_i)\|^2}{\|\vec{x}_i\|^2}$$

where $\{\vec{x}_i\}_{i=1}^n$ are the data points, j is the level in the GMRA, ranging from 0 to J (dependent on the data set), \mathcal{A} is the proposed Algorithm, and M is a fixed random (with respect to Haar measure) orthogonal projection with range of dimension $(d_j \cdot m) \wedge D$, where the “oversampling factor” $m = 1, 2, 4, 16$, and the “intrinsic dimension” $d_j = \max_k \dim(\text{range}(\mathbb{P}_{j,k}))$. Therefore, d_j is the dimension of the manifold (2 and 9, respectively) for \mathcal{M}_1 and \mathcal{M}_2 . The dimension parameter, d_j , is adaptively chosen in a scale-dependent way for $\mathcal{M}_3, \mathcal{M}_4, \mathcal{M}_5$ as described in [3], with actual values used in these examples reported in Figure 1.

There we also run SpaRSA [48] (for reasonable choices of the several parameters involved), one of the leading algorithms, among many, for sparse reconstructions. We notice that: (a) for general real world data sets it achieves comparable precision to our algorithm, suggesting that the GMRA dictionaries may be used in the context of standard sparse approximation; (b) for low-dimensional manifold synthetic data sets, which do not curve in many dimensions, it achieves higher accuracy, since the directions of a few tangent planes are sufficient to span a subspace containing the whole manifold.

Finally, in Figure 2 we report running times, for the same data sets as in Figure 1, for our algorithm \mathcal{A} and SpaRSA. These graphs suggest that our algorithm can perform several orders of magnitude faster than SpaRSA. In the examples shown it took a few seconds to run Algorithm 1 on all the points, with SpaRSA taking a significant fraction of a second to run on a single point.

6. CONCLUSION

In this paper we discussed the ability of random projection to embed an intrinsically low d -dimensional submanifold of \mathbb{R}^D , together with a piecewise linear approximation to the submanifold, into $\mathbb{R}^{O(d \log d)}$ in a way which (approximately) preserves the fidelity of the embedded piecewise linear approximation to the embedded manifold. Although any collection of approximating affine spaces suffice, we focussed on the type of multi-scale linear approximations provided by GMRA [3] in particular. It is worth mentioning that the entire Geometric Wavelet Transform (GWT) [3] of a point near a given manifold can also be preserved by the type of random projections discussed herein.

Note that the GWT of a point on a given manifold will always be approximated by the sum of at most Jd vectors (where J is the number of scales in the GWT). So, pessimistically, a random projection needs to preserve all distances in a number of $O(Jd)$ -dimensional subspaces which is bounded above by Lemma 6 in order to approximately preserve the entire geometric wavelet transform of each point on the manifold. Thus, the GWT of each point on a given manifold should be preserved in compressed form by a random linear projection onto a subspace whose dimension, m , satisfies a variant of Theorem 3 with d replaced everywhere by Jd .

REFERENCES

- [1] Bootstrap methods in signal processing. *IEEE Signal Proc. Mag.*, 24(4), 2007.
- [2] Sensing, sampling, and compression. *IEEE Signal Proc. Mag.*, 25(2), 2008.
- [3] W. Allard, G. Chen, and M. Maggioni. Multiscale geometric methods for data sets II: Geometric wavelets. *Appl. Comp. Harm. Anal.*, accepted, May 2011.
- [4] A. Andoni and P. Indyk. Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions. *Communications of the ACM*, 51(1):419 – 441, 2008.
- [5] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [6] R. Baraniuk and M. Wakin. Random Projections of Smooth Manifolds. *Foundations of Computational Mathematics*, 9(1):51 – 77, 2009.
- [7] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.
- [8] A. Beygelzimer, S. Kakade, and J. Langford. Cover trees for nearest neighbor. In *Proceedings of the 23rd international conference on Machine learning*, pages 97–104. ACM, 2006.




FIGURE 1. From left to right: data sets \mathcal{M}_1 to \mathcal{M}_5 ; from top to bottom: adding noise with increasing variance as above. In each plot, having the horizontal axis representing the scale j and the vertical axis the relative mean square errors $\text{reLMSE}(\mathcal{A}, \mathcal{M}, j)$ as defined in (10), we vary the oversampling parameter $m = 2, 4, 16$, and we also plot (dashed black) the reLMSE_j defined again (10). We report the average result of 10 draws of the random matrix \mathcal{M} , and in dashed lines the standard deviation bands around the mean. Note: the maximum square error is about 10 times larger than the mean square error in all cases (not shown). For \mathcal{M}_3 we have $(d_j)_j = (3, 3, 3, 3, 9, 37, 45, 45)$, and for \mathcal{M}_4 we have $d_j = (9, 7, 7, 7, 6, 6, 9, 36, 61, 66, 66)$, and for \mathcal{M}_5 and $d_j = (83, 51, 33, 21, 43, 50)$. We also run SpARSA [48] (see comments in the text).

- [9] T. Blumensath and M. E. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3):265 – 274, 2009.
- [10] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52:489–509, 2006.
- [11] E. Candes, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [12] E. Candes and T. Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005.
- [13] E. Candes and T. Tao. Near optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. on Information Theory*, 2006.
- [14] E. J. Cands, Y. C. Eldar, D. Needell, and P. Randall. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 31(1):59–73, 2010.
- [15] G. Chen, A. Little, M. Maggioni, and L. Rosasco. Some recent advances in the geometric analysis of point clouds in high dimensions. *Wavelets and Multiscale Analysis: Theory and Applications*, Springer, March 2011.
- [16] G. Chen and M. Maggioni. Multiscale geometric and spectral analysis of plane arrangements. In *Proc. CVPR*, 2011. to appear.
- [17] M. Chen, J. Silva, J. Paisley, C. Wang, D. Dunson, and L. Carin. Compressive Sensing on Manifolds Using a Nonparametric Mixture of Factor Analyzers: Algorithm and Performance Bounds. *IEEE Trans. Signal Processing*, pages 6140–6155, Dec. 2010.
- [18] K. L. Clarkson. Tighter bounds for random projections of manifolds. In *SoCG '08: Proceedings of the Twenty-Fourth Annual Symposium on Computational Geometry*, 2008.
- [19] R. Coifman and M. Maggioni. Multiscale data analysis with diffusion wavelets. *Proc. SIAM Bioinf. Workshop, Minneapolis*, April 2007.
- [20] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker. Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps. *PNAS*, 102(21):7426–7431, 2005.
- [21] S. Dasgupta and A. Gupta. An elementary proof of a theorem of johnson and lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2003.

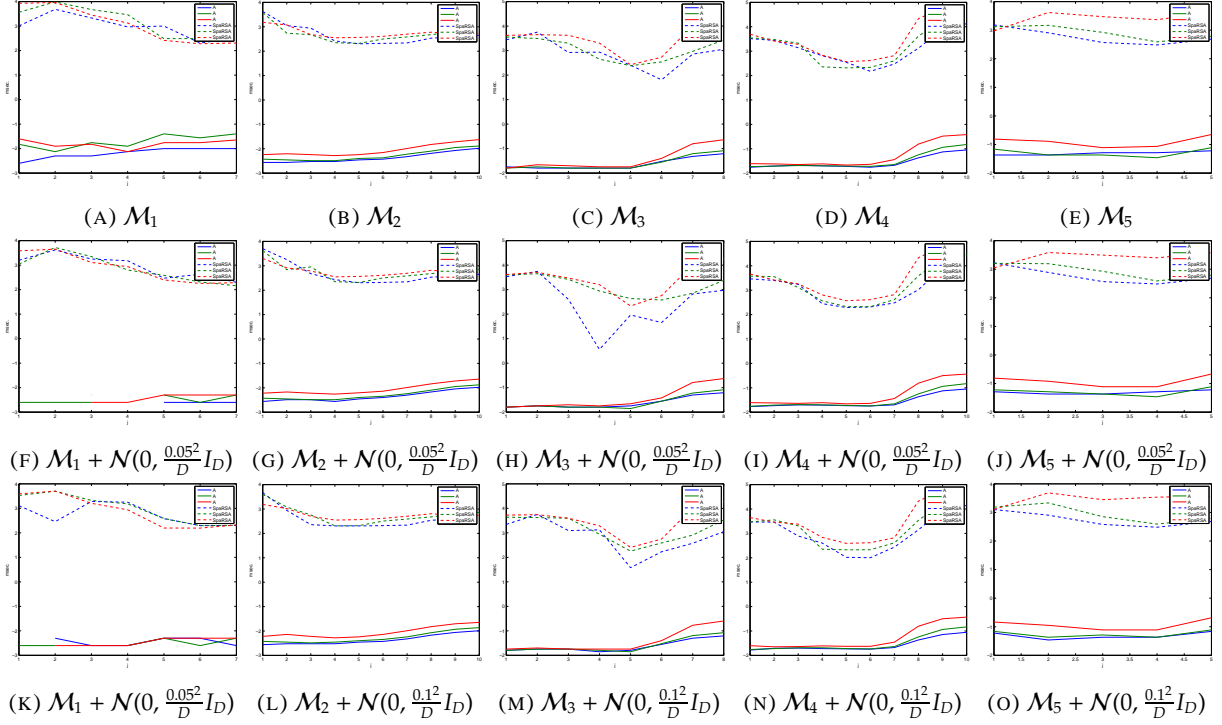


FIGURE 2. From left to right: data sets \mathcal{M}_1 to \mathcal{M}_5 ; from top to bottom: adding noise with increasing variance as above. In each plot, having the horizontal axis representing the scale j and the vertical axis the time (in milliseconds) per point x needed to compute either $\mathcal{A}(x)$ or $\text{SpaRSA}(x)$ [48] (with the matrices involved in the algorithm, and their transposed, precomputed). Because of time involved, we ran SpaRSA only on 50 randomly chosen points rather than all the points.

- [22] D. Donoho. Compressed Sensing. *IEEE Trans. on Information Theory*, 52:1289–1306, 2006.
- [23] D. Donoho and M. Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization. *Proceedings of the National Academy of Sciences*, 100(5):2197–2202, 2003.
- [24] D. L. Donoho and C. Grimes. Hessian eigenmaps: new locally linear embedding techniques for high-dimensional data. *Proc. Nat. Acad. Sciences*, pages 5591–5596, March 2003.
- [25] D. N. E. J. Cands, Y. Eldar and P. Randall. Compressed sensing with coherent and redundant dictionaries. *Appl. Comp. Harm. Anal.*, 31(1):59–73.
- [26] H. Federer. Curvature measures. *Trans. Am. Math. Soc.*, 93(3):418 – 491, 1959.
- [27] P. Frankl and H. Maehara. The johnson-lindenstrauss lemma and the sphericity of some graphs. *Journal of Combinatorial Theory, Series B*, 44(3):355–362, 1988.
- [28] A. C. Gilbert, M. Muthukrishnan, and M. J. Strauss. Approximation of functions over redundant dictionaries using coherence. In *Proc. of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, Jan. 2003.
- [29] G. Haro, G. Randall, and G. Sapiro. Translated Poisson mixture model for stratification learning. *Int. J. Comput. Vision*, 80(3):358–374, 2008.
- [30] W. Hong, J. Wright, K. Huang, and Y. Ma. A multi-scale hybrid linear model for lossy image representation. In *ICCV*, pages 764–771, 2005.
- [31] P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613. ACM, 1998.
- [32] W. B. Johnson and J. Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemp. Math*, 26:189–206, 1984.
- [33] F. Krahmer and R. Ward. New and improved johnson-lindenstrauss embeddings via the restricted isometry property. *SIAM J. Math. Anal.*, 43(3):1269–1281, 2011.
- [34] Y. Ma, A. Y. Yang, H. Derksen, and R. Fossum. Estimation of subspace arrangements with applications in modeling and segmenting mixed data. *SIAM Review*, 50(3):413–458, 2008.
- [35] S. Mallat. A wavelet tour of signal processing. *Academic Press*, 2003.
- [36] D. Needell and J. Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26(3):301–321, 2009.
- [37] D. Needell and R. Vershynin. Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit. *Foundations of Computational Mathematics*, 9:317–334, 2009.

- [38] D. Needell and R. Vershynin. Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit. *IEEE Journal of Selected Topics in Signal Processing*, 4(2):310–316, 2010.
- [39] P. Niyogi, S. Smale, and S. Weinberger. Finding the Homology of Submanifolds with High Confidence from Random Samples. *Discrete Comput Geom*, 39:419 – 441, 2008.
- [40] H. Rauhut. Compressive sensing and structured random matrices. *Theoretical Foundations and Numerical Methods for Sparse Recovery*, 9:1–92, 2010.
- [41] H. Rauhut, K. Schnass, and P. Vandergheynst. Compressed sensing and redundant dictionaries. *IEEE Trans. Inform. Theory*, 54(5):22102219, 2008.
- [42] Y. Sugaya and K. Kanatani. Multi-stage unsupervised learning for multi-body motion segmentation. *IEICE Transactions on Information and Systems*, E87-D(7):1935–1942, 2004.
- [43] J. B. Tenenbaum, V. D. Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- [44] J. Tropp and A. Gilbert. Signal recovery from partial information via orthogonal matching pursuit. *IEEE Trans. Info. Theory*, 53(12):4655–4666, Dec. 2007.
- [45] R. Vidal. Subspace clustering. *IEEE Signal Processing Magazine*, 28(2):52–68, 2011.
- [46] M. B. Wakin. Manifold-Based Signal Recovery and Parameter Estimation from Compressive Measurements. *Preprint*, 2008.
- [47] M. B. Wakin, D. L. Donoho, H. Choi, and R. G. Baraniuk. The multiscale structure of non-differentiable image manifolds. *Proc. Wavelets XI*, SPIE Optics and Photonics, San Diego, California, August, 2005.
- [48] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo. Sparse reconstruction by separable approximation. *IEEE Transactions on Signal Processing*, 57(7):2479–2493, 2009.
- [49] Z. Zhang and H. Zha. Principal manifolds and nonlinear dimension reduction via local tangent space alignment. *SIAM Journal of Scientific Computing*, 26:313–338, 2002.

APPENDIX A. PROOF OF LEMMA 8

To prove this lemma we will modify the proof of Theorem 3.1 in [6]. The proof of Theorem 3.1 proceeds in two steps. First, a finite set, $B \subset \mathbb{R}^D$, of points on/near the given manifold \mathcal{M} is defined. The main body of the proof then consists of demonstrating that any $m \times D$ matrix, M' , which embeds B into \mathbb{R}^m with $\Theta(\epsilon)$ -distortion will also satisfy

$$(1 - \epsilon) \|\vec{x} - \vec{y}\| \leq \|M' \vec{x} - M' \vec{y}\| \leq (1 + \epsilon) \|\vec{x} - \vec{y}\|$$

for all $\vec{x}, \vec{y} \in \mathcal{M}$. Our proof will proceed along a similar path. We will begin by first defining a modified version of the set, B , considered in [6]. We will call this set \tilde{B} . Then, we will prove that any $m \times D$ matrix which which embeds \tilde{B} into \mathbb{R}^m with $\Theta(\epsilon)$ -distortion will also satisfy item (a) of Assumption Set 2 in Section 2.3.

Let $d_{\mathcal{M}}(\vec{x}, \vec{y})$ denote the geodesic distance between $\vec{x}, \vec{y} \in \mathcal{M}$. Furthermore, let $\mathbf{Tan}_{\vec{x}}$ denote the d -dimensional tangent space to \mathcal{M} at each $\vec{x} \in \mathcal{M}$. Finally, let

$$\mathcal{B}_{\mathcal{M}, \delta}(\vec{x}) = \{\vec{y} \in \mathcal{M} \mid d_{\mathcal{M}}(\vec{x}, \vec{y}) \leq \delta\}$$

for each $\delta \in \mathbb{R}^+$ and $\vec{x} \in \mathcal{M}$.

We are now ready to construct $B \subset \mathbb{R}^D$ as per [6] as follows: Set $T = O\left(\frac{\epsilon^2}{D} \cdot \min\{1, \text{reach}(\mathcal{M})\}\right)$ and, for each $\vec{x} \in \mathcal{M}$, let $Q_2(\vec{x}) \subset \mathbf{Tan}_{\vec{x}}$ denote a minimal $\Theta(\epsilon \cdot T / \sqrt{D})$ -cover of the d -dimensional Euclidean ball of radius T centered at $\vec{0} \in \mathbf{Tan}_{\vec{x}}$. Next, choose $A \subset \mathcal{M}$ to be a minimal finite cover of \mathcal{M} satisfying

$$\min_{\vec{a} \in A} d_{\mathcal{M}}(\vec{a}, \vec{x}) \leq T,$$

for all $\vec{x} \in \mathcal{M}$. Then,

$$B := \bigcup_{\vec{a} \in A} \{\vec{a}\} \cup (\vec{a} + Q_2(\vec{a})).$$

In the next paragraph we will define our modified set, $\tilde{B} \subset \mathbb{R}^D$, which is a superset of the set B defined above.

Fix $j \in [J]$ and $k \in [K_j]$. For each $\vec{a} \in A$ above, let $\vec{a}_{j,k} \in \mathcal{B}_{\mathcal{M}, T}(\vec{a})$ be such that

$$\|\vec{a}_{j,k} - \vec{c}_{j,k}\| \leq \|\vec{y} - \vec{c}_{j,k}\| \quad \forall \vec{y} \in \mathcal{B}_{\mathcal{M}, T}(\vec{a}).$$

Let $A_{j,k} = \{\vec{a}_{j,k} \mid \vec{a} \in A\}$. Furthermore, denote the $(d+1)$ -dimensional vector space spanned by $\mathbf{Tan}_{\vec{a}_{j,k}} \cup \{\vec{c}_{j,k} - \vec{a}_{j,k}\}$ by $\mathbf{Tan}_{\vec{a}_{j,k}}$, and then let $Q_{j,k}(\vec{a}) \subset \mathbf{Tan}_{\vec{a}_{j,k}}$ be a minimal $\Theta(\epsilon \cdot T / \sqrt{D})$ -cover of the $(d+1)$ -dimensional Euclidean ball of radius T centered at $\vec{0}$. To finish, define

$$B_{j,k} := \bigcup_{\vec{a} \in A} \{\vec{a}_{j,k}\} \cup (\vec{a}_{j,k} + Q_{j,k}(\vec{a}))$$

and then set

$$\tilde{B} := \left(\bigcup_{j \in [J], k \in [K_j]} B_{j,k} \cup \{\vec{c}_{j,k}\} \right) \cup B \cup Q,$$

where $Q \subset \mathbb{R}^D$ is as defined in Lemma 7.

Note that $|\tilde{B}|$ will be bounded above by

$$(J+1) \cdot K_J \left(1 + \max_{j \in [J], k \in [K_j]} |B_{j,k}| \right) + |B| + |Q|.$$

Applying Lemma 7 to bound $|Q|$, Lemma 6 to bound K_J , and appealing to Section 3.2.5 of [6] to bound $|B|$, the previous line reveals that

$$(11) \quad |\tilde{B}| \ll 2^{O(J \cdot d)} \cdot V \cdot \left(\frac{d}{\min\{1, C_1\}} \right)^{O(d)} \left(\max_{j \in [J], k \in [K_j]} |B_{j,k}| \right) + V \left(\frac{D}{\epsilon \cdot \min\{1, \text{reach}(\mathcal{M})\}} \right)^{O(d)}.$$

We now finish bounding the cardinality of \tilde{B} by noting that $|B_{j,k}|$ will always be bounded above by the upper bounds for $|B|$ in Section 3.2.5 of [6] after every occurrence of $K = d$ is replaced with $d + 1$.⁵ The stated upper bound on $|\tilde{B}|$ follows.

We will now complete the second portion of our proof by demonstrating that a sufficiently precise linear embedding of \tilde{B} will satisfy item (a) of Assumption Set 2. First, since $B \subset \tilde{B}$, Theorem 3.1 in [6] guarantees that a low-distortion embedding of \tilde{B} will preserve all pairwise distances between points on the manifold \mathcal{M} . Furthermore, any embedding of \tilde{B} will also embed all $\vec{c}_{j,k}$ -vectors since they form a proper subset of \tilde{B} . Hence, it suffices for us to show that a sufficiently precise linear embedding of \tilde{B} will (approximately) preserve the distance from each $\vec{c}_{j,k}$ -vector to all points on the manifold \mathcal{M} .

Fix $j \in [J]$, $k \in [K_j]$, and $\vec{x} \in \mathcal{M}$. Let $\vec{a}' \in A$ be the closest element of A to \vec{x} ,

$$\vec{a}' = \arg \min_{\vec{a} \in A} d_{\mathcal{M}}(\vec{a}, \vec{x}).$$

Finally, let $\vec{x}'_{j,k}$ denote the projection of \vec{x} onto the $(d+1)$ -dimensional affine subspace $\vec{a}'_{j,k} + \mathbf{Tan}_{\vec{a}'_{j,k}}$. By considering the Taylor series expansion of the unit speed parameterization of the geodesic path from $\vec{a}'_{j,k}$ to \vec{x} on \mathcal{M} , we find that

$$\vec{x} = \vec{x}'_{j,k} + \vec{r}, \text{ where } \|\vec{r}\| = O\left(\frac{d_{\mathcal{M}}^2(\vec{x}, \vec{a}'_{j,k})}{\text{reach}(\mathcal{M})}\right).$$

In fact, the magnitude of the remainder, \vec{r} , is also $O\left(\|\vec{x} - \vec{a}'_{j,k}\|^2\right)$ since $T < \text{reach}(\mathcal{M})/2$ (see Corollary 2.1 in [6]).

Furthermore, the definition of $\vec{a}'_{j,k} \in \mathcal{M}$ implies that $\|\vec{x} - \vec{a}'_{j,k}\| = O(\|\vec{x} - \vec{c}_{j,k}\|)$.

Continuing with the proof, suppose that an $m \times D$ matrix, M' , embeds \tilde{B} into \mathbb{R}^m with $\Theta(\epsilon)$ -distortion. A trivial variant of Lemma 5 then implies that

$$\begin{aligned} \|M'\vec{x} - M'\vec{c}_{j,k}\| &\leq \|M'\vec{x} - M'\vec{x}'_{j,k}\| + \|M'\vec{x}'_{j,k} - M'\vec{c}_{j,k}\| \leq \|M'\vec{r}\| + (1 + \Theta(\epsilon)) \|\vec{x}'_{j,k} - \vec{c}_{j,k}\| \\ &\leq (1 + \Theta(\epsilon)) (\|\vec{x} - \vec{c}_{j,k}\| + \|\vec{r}\|) + \|M'\vec{r}\| \leq (1 + \Theta(\epsilon)) \|\vec{x} - \vec{c}_{j,k}\| + \|M'\vec{r}\| + O\left(\|\vec{x} - \vec{a}'_{j,k}\|^2\right) \end{aligned}$$

since $Q_{j,k}(\vec{a}') \subset \mathbf{Tan}_{\vec{a}'_{j,k}}$ is a proper subset of \tilde{B} , and $(\vec{x}'_{j,k} - \vec{c}_{j,k}) \in \mathbf{Tan}_{\vec{a}'_{j,k}}$. In addition, the fact that $Q \subset \tilde{B}$ together with Lemma 7 guarantees that M' will have the $\text{RIP}(D, d, \Theta(\epsilon))$. This fact combined with the Hölder inequality finally reveals that

$$\begin{aligned} \|M'\vec{x} - M'\vec{c}_{j,k}\| &\leq (1 + \Theta(\epsilon)) \|\vec{x} - \vec{c}_{j,k}\| + O\left(\sqrt{\frac{D}{d}} \cdot \|\vec{x} - \vec{a}'_{j,k}\|^2\right) \leq \left(1 + \Theta(\epsilon) + O\left(\sqrt{\frac{D}{d}} \cdot T\right)\right) \|\vec{x} - \vec{c}_{j,k}\| \\ &\leq (1 + O(\epsilon)) \|\vec{x} - \vec{c}_{j,k}\|. \end{aligned}$$

The lower bound for $\|M'\vec{x} - M'\vec{c}_{j,k}\|$ is established in an analogous fashion. We have the stated theorem.

⁵Intuitively, we are increasing the effective intrinsic dimensionality of \mathcal{M} from d to $d + 1$ in the process of creating our $B_{j,k}$ -subsets.

APPENDIX B. PROOF OF LEMMA 9

The proof of this Lemma borrows heavily from the proof of Lemma 8. Set $T = O\left(\frac{2^{-J}\epsilon^2}{D} \cdot \min\{1, \text{reach}(\mathcal{M})\}\right)$. We will begin by defining the set $B' \subset \mathbb{R}^D$. Let $A \subset \mathcal{M}$, $B \subset \mathbb{R}^D$, and $A_{j,k} = \{\vec{a}_{j,k} \mid \vec{a} \in A\} \subset \mathcal{M}$ for each $j \in [J], k \in [K_j]$ be defined as in Appendix A above (except now using the smaller value of T from the second sentence of this appendix). Let $\widetilde{\mathbf{Tan}}_{\vec{a}_{j,k}}$ denote the $(2d+1)$ -dimensional vector space spanned by

$$\mathbf{Tan}_{\vec{a}_{j,k}} \bigcup \{\vec{c}_{j,k} - \vec{a}_{j,k}\} \bigcup \{\Phi_{j,k}^T \Phi_{j,k} \vec{y} \mid \vec{y} \in \mathbb{R}^D\}$$

for each $\vec{a}_{j,k} \in A_{j,k}$. Furthermore, for each $\vec{a}_{j,k} \in A_{j,k}$, let $Q'_{j,k}(\vec{a}) \subset \widetilde{\mathbf{Tan}}_{\vec{a}_{j,k}}$ be a minimal $\Theta(\epsilon \cdot T / \sqrt{D})$ -cover of the $(2d+1)$ -dimensional Euclidean ball of radius T centered at $\vec{0} \in \widetilde{\mathbf{Tan}}_{\vec{a}_{j,k}}$. To finish, define

$$B'_{j,k} := \bigcup_{\vec{a} \in A} \{\vec{a}_{j,k}\} \cup (\vec{a}_{j,k} + Q'_{j,k}(\vec{a}))$$

for each $j \in [J], k \in [K_j]$, and then set

$$B' := \left(\bigcup_{j \in [J], k \in [K_j]} B'_{j,k} \cup \{\vec{c}_{j,k}\} \right) \cup B \cup Q,$$

where $Q \subset \mathbb{R}^D$ is as defined in Lemma 7. It is not difficult to see that $|B'|$ will be bounded above as per Equation 11 after ϵ is replaced everywhere by $2^{-J}\epsilon$. Simplifying yields the stated upper bound.

We will now complete our proof by demonstrating that a sufficiently precise linear embedding of B' will satisfy item (d) of Assumption Set 2. Fix $j \in [J], k \in [K_j]$, and $\vec{x} \in \mathcal{M}$. Let $\vec{a}' \in A$ be the closest element of A to \vec{x} ,

$$\vec{a}' = \arg \min_{\vec{a} \in A} d_{\mathcal{M}}(\vec{a}, \vec{x}).$$

Finally, let $\vec{x}'_{j,k}$ denote the projection of \vec{x} onto the $(2d+1)$ -dimensional affine subspace $\vec{a}'_{j,k} + \widetilde{\mathbf{Tan}}_{\vec{a}_{j,k}}$. By considering the Taylor series expansion of the unit speed parameterization of the geodesic path from $\vec{a}'_{j,k}$ to \vec{x} on \mathcal{M} , we find that

$$\vec{x} = \vec{x}'_{j,k} + \vec{r}, \text{ where } \|\vec{r}\| = O\left(\frac{d_{\mathcal{M}}^2(\vec{x}, \vec{a}'_{j,k})}{\text{reach}(\mathcal{M})}\right).$$

Furthermore, we recall that the magnitude of the remainder, \vec{r} , is also $O\left(\|\vec{x} - \vec{a}'_{j,k}\|^2\right)$ since T is sufficiently small.

To finish, suppose that an $m \times D$ matrix, M' , embeds B' into \mathbb{R}^m with $\Theta(\epsilon)$ -distortion. A trivial variant of Lemma 5 implies that

$$\begin{aligned} \|M' \vec{x} - M' \mathbb{P}_{j,k}(\vec{x})\| &\leq \|M' \vec{x} - M' \vec{x}'_{j,k}\| + \|M' \vec{x}'_{j,k} - M' \mathbb{P}_{j,k}(\vec{x})\| \leq \|M' \vec{r}\| + (1 + \Theta(\epsilon)) \|\vec{x}'_{j,k} - \mathbb{P}_{j,k}(\vec{x})\| \\ &\leq (1 + \Theta(\epsilon)) (\|\vec{x} - \mathbb{P}_{j,k}(\vec{x})\| + \|\vec{r}\|) + \|M' \vec{r}\| \\ &\leq (1 + \Theta(\epsilon)) \|\vec{x} - \mathbb{P}_{j,k}(\vec{x})\| + \|M' \vec{r}\| + O\left(\|\vec{x} - \vec{a}'_{j,k}\|^2\right) \end{aligned}$$

since $Q'_{j,k}(\vec{a}') \subset \widetilde{\mathbf{Tan}}_{\vec{a}_{j,k}}$ is a subset of B' , and $(\vec{x}'_{j,k} - \mathbb{P}_{j,k}(\vec{x})) \in \widetilde{\mathbf{Tan}}_{\vec{a}_{j,k}}$. In addition, the fact that $Q \subset B'$ together with Lemma 7 guarantees that M' will have the $\text{RIP}(D, d, \Theta(\epsilon))$. This fact combined with the Hölder inequality reveals that

$$\begin{aligned} \|M' \vec{x} - M' \mathbb{P}_{j,k}(\vec{x})\| &\leq (1 + \Theta(\epsilon)) \|\vec{x} - \mathbb{P}_{j,k}(\vec{x})\| + O\left(\sqrt{\frac{D}{d}} \cdot \|\vec{x} - \vec{a}'_{j,k}\|^2\right) \leq (1 + \Theta(\epsilon)) \|\vec{x} - \mathbb{P}_{j,k}(\vec{x})\| + O\left(\sqrt{\frac{D}{d}} \cdot T^2\right) \\ &\leq (1 + \Theta(\epsilon)) \|\vec{x} - \mathbb{P}_{j,k}(\vec{x})\| + 2^{-J} \end{aligned}$$

whenever T is weighted by a sufficiently small (universal) constant. The lower bound for $\|M' \vec{x} - M' \mathbb{P}_{j,k}(\vec{x})\|$ is established in an analogous fashion.